A Formal Approach to Cyber-Physical Attacks

Ruggero Lanotte
Dipartimento di Scienza e Alta Tecnologia
Università dell'Insubria, Como, Italy

Massimo Merro and Riccardo Muradore
Dipartimento di Informatica
Università degli Studi di Verona, Italy

Luca Viganò
Department of Informatics
King's College London, UK

arXiv:1611.01377v1 [cs.CR] 4 Nov 2016

***Abstract*—Cyber-Physical Systems (CPSs) are integrations of networking and distributed computing systems with physical processes that monitor and control entities in a physical environment, with feedback loops where physical processes affect computations and vice versa. In this paper, we apply formal methods to lay and streamline theoretical foundations to reason about CPSs and cyber-physical attacks. We focus on a formal treatment of both integrity and DoS attacks to sensors and actuators of CPSs, paying particular attention to the timing aspects of these attacks. Our contributions are threefold: (1) we define a hybrid process calculus to model both CPSs and cyber-physical attacks; (2) we define a threat model of cyber-physical attacks and provide the means to assess attack tolerance/vulnerability with respect to a given attack; (3) we formalise how to estimate the impact of a successful attack on a CPS and investigate possible quantifications of the success chances of an attack. We illustrate definitions and results by means of a non-trivial engineering application.**

1. Introduction

Context and motivation. *Critical system (or infrastructure)* is a term used to describe key assets and sectors that are essential for the functioning of a society, such as services and facilities for telecommunication, finances, government and administration, public health, transportation, energy and water supply, smart grids/homes/cities, manufacturing, etc. The informatisation of critical systems has been growing steadily thanks to the progress in the development and usage of *Cyber-Physical Systems* (CPSs), which are integrations of networking and distributed computing systems with physical processes that monitor and control entities in the physical environment, with feedback loops where physical processes affect computations and vice versa. For example, in real-time control systems, a hierarchy of *sensors*, *actuators* and *control processing components* are connected to control stations. Different kinds of CPSs include supervisory control and data acquisition (SCADA), programmable logic controllers (PLC) and distributed control systems.

The need for establishing the security of critical systems, and thus of CPSs, is well understood, especially since system failure would be extremely costly and threaten not only the system's environment, but also the existence of the organisation that operates the system or, ultimately, even human life. However, this need is often not addressed in current industrial practice because of the major challenges that are posed by the heterogeneous and distributed nature of the systems and their interaction with the physical world and, sometimes, with the human users.

As a consequence, the last years have seen a dramatic increase in the number of attacks to the security of cyber-physical and critical systems, e.g., manipulating sensor readings and, in general, influencing physical processes to bring the system into a state desired by the attacker. Many (in)famous examples have been so impressive to make the international news, e.g.: the Stuxnet worm, which reprogrammed PLCs of nuclear centrifuges in Iran [6]; the attack on a sewage treatment facility in Queensland, Australia, which manipulated the SCADA system to release raw sewage into local rivers and parks [23]; and the SQL Slammer worm, which affected a private computer network at a nuclear power plant in Ohio [15].

As observed in [13], until now the primary focus of academia and industry has been on securing the communication infrastructure and improving robustness and fault-tolerance of control systems, and there is a large body of literature on how to adapt existing IT-security methods to the characteristic features of the control domain [22]. However, a common feature of most of the well-known attacks, including those mentioned above, is that they succeeded because the attacks were able to manipulate the systems' sensors, actuators, PLCs and SCADA components irregardless of the cryptographic security measures implemented in the systems to protect the underlying communication infrastructure.

In other words, such CPS attacks were not only due to the cryptographic communication protocols employed, and we cannot address these security issues only by resorting to approaches and tools that have been specifically developed for cryptographic protocol analysis and network security in general. More specifically, state-of-the-art technologies for network security modelling and analysis do not provide satisfactory support for the discovery of important vulnerabilities and associated exploits that are already plaguing informatised CPSs, and thus severely affect the further development and adoption of these systems. A number of approaches and tools have proven to be quite successful for the formal analysis of Internet security protocols, services and applications, validating realistic case studies and uncovering a number of vulnerabilities even in widely used solutions. CPSs, however,

exhibit idiosyncrasies that make them unsuitable to the direct application of such approaches and tools. To carry out automated security analysis of such critical systems, proper extensions of the modelling and reasoning techniques implemented in existing tools are needed to be able to express and reason about the wide spectrum of the systems' features.

As stated in [9], the concern for consequences at the physical level puts *CPS security* apart from standard *information security*, and demands for *ad hoc* solutions to properly address such novel research challenges. The works that have taken up these challenges range from proposals of different notions of cyber-physical security and attacks (e.g., [3], [9], [13], to name a few) to pioneering extensions to CPS security of standard formal approaches (e.g., [3], [4], [27]). However, to the best of our knowledge, a systematic formal approach to cyber-physical attacks is still to be fully developed.

Background. The dynamic behaviour of the *physical plant* of a CPS is usually represented by means of a discrete-time state-space model [28]. Typically, it can be modelled by equations of the form

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k + Cw_k \\ y_k &= Dx_k + Ee_k \end{aligned}$$

where $x_k \in \mathbb{R}^n$ is the current *(physical) state*, $u_k \in \mathbb{R}^m$ is the *input* (i.e., the control actions implemented through actuators) and $y_k \in \mathbb{R}^p$ is the *output* (i.e., the measurements from the sensors). The *uncertainty* $w_k \in \mathbb{R}^n$ and the *measurement error* $e_k \in \mathbb{R}^p$ represent the discrepancies between the model and the real process, due to unmodelled dynamics or disturbances.[1] Here, the *next state* x_{k+1} depends on the current state x_k and the corresponding control actions u_k, at the sampling instant $k \in \mathbb{N}$. The state x_k cannot be directly observed: only its measurements y_k can be observed.

The physical plant is supported by a communication network through which the sensor measurements and actuator data are exchanged with controller(s) and supervisor(s) (e.g., IDSs), which are the *cyber* components (also called *logics*) of a CPS.

Contributions. In this paper, we focus on a formal treatment of both *integrity* and *Denial of Service (DoS)* attacks to *physical devices* (sensors and actuators) of CPSs, considering in particular the *timing aspects* of these attacks. Our overall goal is to apply formal methodologies to lay *theoretical foundations* to reason about and detect attacks to physical devices of CPSs. That is, to provide an essential stepping stone for formal and automated analysis techniques for checking the security of CPSs (rather than for providing defence techniques).

Our contributions are threefold. As a first contribution, we define a *hybrid process calculus*, called CCPSA, to formally specify both CPSs and cyber-physical attacks. In CCPSA, CPSs are represented as terms of the form

$$E \bowtie_{\mathcal{S}} P$$

1. A, B, C, D and E are matrices modelling the dynamics of the physical system.

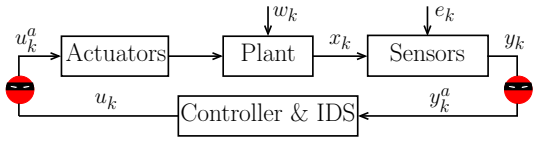

Figure 1. Our threat model for CPSs

where E denotes the *physical plant* (also called environment) of the system, containing information on state variables, actuators, sensors, evolution law, etc., and P represents the *cyber component* of the system that governs access to sensors and actuators, communication with other cyber components, and detection of potentially dangerous physical states. Thus, channels are used for logical interactions between cyber components, whereas sensors and actuators make possible the interaction between cyber and physical components. Finally, $\mathcal{S}$ denotes a set of sensors and actuators of the CPS that are "secured" using some cryptographic communication. Those devices are thus protected from attacks.

CCPSA adopts a *discrete notion of time* [10] and is equipped with two trace-based *behavioural equalities*: a *trace equivalence* semantics and a variant of it that takes into account discrepancies of execution traces at specific time instants.

As a second contribution, we formalise a *threat model* that specifies attacks that drive a CPS into an undesired state [24]. Cyber-physical attacks typically tamper with both the physical and the cyber layer. In our threat model, depicted in Figure 1, communication is observable but cannot be manipulated by the attacker, who instead may manipulate *unprotected* physical devices, which is our focus. As depicted in Figure 1, our attacks may affect directly the sensor measurements or the controller commands.

- *Attacks on sensors* consist of reading and eventually replacing y_k (the sensor measurements) with y_k^a.
- *Attacks on actuators* consist of reading, eavesdropping and eventually replacing the controller commands u_k with u_k^a, affecting directly the actions the actuators may execute.

We group attacks into classes of the form $\mathcal{I}[m, n]$, which describe both the malicious activity $\mathcal{I}$ on the physical devices and the *timing parameters* m and n of the attack: begin and end of the attack. As observed in [13], timing is a critical issue in CPSs because the physical state of a system changes continuously over time, and as the system evolves in time, some states might be more vulnerable to attacks than others. For example, an attack launched when the target state variable reaches a local maximum (or minimum) may have a great impact on the whole system behaviour [14]. Furthermore, not only the timing of the attack but also the duration of the attack is an important parameter to be taken into consideration in order to achieve a successful attack. For example, it may take minutes for a chemical reactor to rupture [25], hours to heat a tank of water or burn out a motor, and days to destroy centrifuges [6].

Then, we use the trace-based behavioural semantics

of CCPSA to formally define when a CPS *tolerates* a cyber-physical attack of a certain class. Intuitively, this happens when the observational behaviour of the CPS under attack remains unchanged when compared with its genuine behaviour, i.e., when running in isolation. We then use our time-sensitive trace equivalence to formally define when a CPS is *vulnerable* to an attack A of a certain class $\mathcal{I}[m, n]$. Intuitively, this happens if there is a time interval $[m', n']$ in which the observational behaviour of the CPS under attack deviates from the correct one. Thus, if a system is vulnerable to an attack of class $\mathcal{I}[m, n]$, then the time interval $[m', n']$ allow us to model: (i) *stealthy attacks* (terminating before producing observable effects, i.e., $n < m'$); (ii) *transitory attacks* (when at some point the system recovers its correct behaviour, i.e., $n' < \infty$); (iii) *permanent attacks* (when the system under attack will never recover its correct behaviour, i.e., $n' = \infty$). Moreover, these parameters may be useful to estimate the time available for IDSs to react before attacks achieve their goal (this time is $m' - n$).

As a third contribution, we formalise how to estimate the *impact of a successful attack on a CPS* and investigate possible *quantifications of the chances* for an attack of being successful when attacking a CPS. We provide a *metric* to estimate the deviation of the system under attack with respect to expected behaviour, according to its evolution law and the uncertainty of the model. This will allow the attacker to reason about the costs of an attack as well as allow the engineers to reason about the cost of stopping a plant under attack.

We illustrate our definitions and results by means of a *running example* taken from an engineering application. While we have kept the example simple, it is actually far from trivial and designed to describe a wide number of attacks.

All the attacks exhibited in the paper have been formally proved, but proofs are sketched or omitted altogether; details are given in the appendix. Furthermore, the behaviour of our running example and of most of the cyber-physical attacks appearing in the paper have been simulated in MATLAB.

Organisation. In § 2, we give syntax and semantics of CCPSA. In § 3, we define cyber-physical attacks and attack tolerance/vulnerability. In § 4, we estimate the impact of a successful attack on a CPS and investigate possible quantifications of the success chances of an attack. In § 5, we draw conclusions and discuss related and future work.

2. A Calculus of Cyber-Physical Systems and Attacks

In this section, we introduce our *Calculus of Cyber-Physical Systems and Attacks* CCPSA. We begin by providing our formalisation of cyber-physical systems equipped with a labelled transition semantics and a trace-based behavioural semantics.

2.1. A formalisation of cyber-physical systems

Let us start with some preliminary notations. We use

- $x, x_k \in \mathcal{X}$ for *state variables*,
- $c, d \in \mathcal{C}$ for *communication channels*,
- $a, a_k \in \mathcal{A}$ for *actuator devices*,
- $s, s_k \in \mathcal{S}$ for *sensors devices*, and
- p, q for both sensors and actuators (generically called *physical devices*).

Values, ranged over by $v, v' \in \mathcal{V}$, are built from basic values, such as Booleans, integers and real numbers. *Actuator names* are metavariables for actuator devices like $valve$, $light$, etc. Similarly, *sensor names* are metavariables for sensor devices, e.g., a sensor $thermometer$ that measures, with a given precision, a state variable called $temperature$.

Given a generic set of names $\mathcal{N}$, we write $\mathbb{R}^{\mathcal{N}}$ to denote the set of functions $[\mathcal{N} \to \mathbb{R}]$ assigning a real value to each name in $\mathcal{N}$. For $\xi \in \mathbb{R}^{\mathcal{N}}$, $n \in \mathcal{N}$ and $v \in \mathbb{R}$, we write $\xi[n \mapsto v]$ to denote the function $\psi \in \mathbb{R}^{\mathcal{N}}$ such that $\psi(m) = \xi(m)$, for any $m \neq n$, and $\psi(n) = v$.

Definition 1 (Cyber-physical system)**.** *In CCPSA, a* cyber-physical system *consists of two components: a* physical environment E *that encloses all physical aspects of a system (state variables, physical devices, evolution law, etc) and a* cyber component, represented as a concurrent process P *that interacts with the physical devices (sensors and actuators) of the system, and can communicate, via channels, with other processes of the same CPS or with processes of other CPSs.*

Given a set $\mathcal{S}$ of secured physical devices *of E, we write $E \bowtie_{\mathcal{S}} P$ to denote the resulting CPS, and use M and N to range over CPSs. We write $E \bowtie P$ when $\mathcal{S} = \emptyset$.*

Intuitively, the "secured" devices in $\mathcal{S}$ are accessed in a protected way (using some cryptographic communication) and hence they cannot be attacked.

Let us now define physical environments E and concurrent processes P in order to formalise our proposal for modelling (and reasoning about) CPSs and their components.

Definition 2 (Physical environment)**.** *Let $\hat{\mathcal{X}} \subseteq \mathcal{X}$ be a set of state variables, $\hat{\mathcal{A}} \subseteq \mathcal{A}$ be a set of actuators, and $\hat{\mathcal{S}} \subseteq \mathcal{S}$ be a set of sensors. A* physical environment E *is 7-tuple $\langle \xi_{\mathrm{x}}, \xi_{\mathrm{u}}, \xi_{\mathrm{w}}, evol, \xi_{\mathrm{e}}, meas, inv \rangle$, where:*

- $\xi_{\mathrm{x}} \in \mathbb{R}^{\hat{\mathcal{X}}}$ *is the* state function*,*
- $\xi_{\mathrm{u}} \in \mathbb{R}^{\hat{\mathcal{A}}}$ *is the* actuator function*,*
- $\xi_{\mathrm{w}} \in \mathbb{R}^{\hat{\mathcal{X}}}$ *is the* uncertainty function*,*
- $evol : \mathbb{R}^{\hat{\mathcal{X}}} \times \mathbb{R}^{\hat{\mathcal{A}}} \times \mathbb{R}^{\hat{\mathcal{X}}} \to 2^{\mathbb{R}^{\hat{\mathcal{X}}}}$ *is the* evolution map*,*
- $\xi_{\mathrm{e}} \in \mathbb{R}^{\hat{\mathcal{S}}}$ *is the* sensor-error function*,*
- $meas : \mathbb{R}^{\hat{\mathcal{X}}} \times \mathbb{R}^{\hat{\mathcal{S}}} \to 2^{\mathbb{R}^{\hat{\mathcal{S}}}}$ *is the* measurement map*,*
- $inv : \mathbb{R}^{\hat{\mathcal{X}}} \to \{\mathsf{true}, \mathsf{false}\}$ *is the* invariant function*.*

All the functions defining an environment are total functions*.*

The *state function* ξ_{x} returns the current value (in $\mathbb{R}$) associated to each state variable of the system. The *actuator function* ξ_{u} returns the current value associated to each actuator. The *uncertainty function* ξ_{w} returns the uncertainty associated to each state variable. Thus, given a state variable

$x \in \hat{\mathcal{X}}$, $\xi_{\mathrm{w}}(x)$ returns the maximum distance between the real value of x and its representation in the model. Later in the paper, we will be interested in comparing the accuracy of two systems. Thus, e.g., for $\xi_{\mathrm{w}}, \xi'_{\mathrm{w}} \in \mathbb{R}^{\hat{\mathcal{X}}}$, we will write $\xi_{\mathrm{w}} \leq \xi'_{\mathrm{w}}$ if $\xi_{\mathrm{w}}(x) \leq \xi'_{\mathrm{w}}(x)$, for any $x \in \hat{\mathcal{X}}$. Both the state function and the actuator function are supposed to change during the evolution of the system, whereas the uncertainty function is supposed to be constant.

Given a state function, an actuator function, and an uncertainty function, the *evolution map* $evol$ returns the set of next *admissible state functions*. This function models the *evolution law* of the physical system, where changes made on actuators may reflect on state variables. Since we assume an uncertainty in our models, the evolution map does not return a single state function but a set of possible state functions. The evolution map is obviously monotone with respect to uncertainty: if $\xi_{\mathrm{w}} \leq \xi'_{\mathrm{w}}$ then $evol(\xi_{\mathrm{x}}, \xi_{\mathrm{u}}, \xi_{\mathrm{w}}) \subseteq evol(\xi_{\mathrm{x}}, \xi_{\mathrm{u}}, \xi'_{\mathrm{w}})$. Note also that, although the uncertainty function is constant, it can be used in the evolution map in an arbitrary way (e.g., it could have a heavier weight when a state variable reaches extreme values).

The *sensor-error function* ξ_{e} returns the maximum error associated to each sensor in $\mathcal{S}$. Again due to the presence of the sensor-error function, the *measurement map* $meas$, given the current state function, returns a set of admissible measurement functions rather than a single one.

Finally, the *invariant function* inv represents the conditions that the state variables must satisfy to allow for the evolution of the system. A CPS whose state variables don't satisfy the invariant is in *deadlock*.

Let us now introduce a running example that will be useful to illustrate the main definitions and results of our approach. We remark that while we have kept the example simple, it is actually far from trivial and designed to describe a wide number of attacks. A more complex example (say, with n sensors and m actuators) wouldn't have been more instructive but just made the paper more dense.

Example 1 (Physical environment of the CPS *Sys*)**.** *Consider a CPS Sys in which the temperature of an engine is maintained within a specific range by means of a cooling system. The physical environment Env of Sys is constituted by: (i) a state variable $temp$ containing the current temperature of the engine; (ii) an actuator $cool$ to turn on/off the cooling system; (iii) a sensor s_{t} (such as a thermometer or a thermocouple) measuring the temperature of the engine; (iv) an uncertainty $\delta = 0.4$ associated to the only variable $temp$; (v) a simple evolution law that increases (respectively, decreases) the value of $temp$ of one degree per time unit if the cooling system is inactive (respectively, active) — the evolution law is obviously affected by the uncertainty δ; (vi) an error $\epsilon = 0.1$ associated to the only sensor s_{t}; (vii) a measurement map to get the values detected by sensor s_{t}, up to its error ϵ; (viii) an invariant function saying that the system gets faulty when the temperature of the engine gets out of the range $[0, 50]$.*

Formally, $Env = \langle \xi_{\mathrm{x}}, \xi_{\mathrm{u}}, \xi_{\mathrm{w}}, evol, \xi_{\mathrm{e}}, meas, inv \rangle$ with:

- $\xi_{\mathrm{x}} \in \mathbb{R}^{\{temp\}}$ *and* $\xi_{\mathrm{x}}(temp) = 0$*;*
- $\xi_{\mathrm{u}} \in \mathbb{R}^{\{cool\}}$ *and* $\xi_{\mathrm{u}}(cool) = \mathsf{off}$*; for the sake of simplicity, we can assume ξ_{u} to be a mapping $\{cool\} \to \{\mathsf{on}, \mathsf{off}\}$ such that $\xi_{\mathrm{u}}(cool) = \mathsf{off}$ if $\xi_{\mathrm{u}}(cool) \geq 0$, and $\xi_{\mathrm{u}}(cool) = \mathsf{on}$ if $\xi_{\mathrm{u}}(cool) < 0$;*
- $\xi_{\mathrm{w}} \in \mathbb{R}^{\{temp\}}$ *and* $\xi_{\mathrm{w}}(temp) = 0.4 = \delta$*;*
- $evol(\xi^i_{\mathrm{x}}, \xi^i_{\mathrm{u}}, \xi_{\mathrm{w}}) = \{\xi : \xi(temp) = \xi^i_{\mathrm{x}}(temp) + heat(\xi^i_{\mathrm{u}}, cool) + \gamma \quad \wedge \quad \gamma \in [-\delta, +\delta]\}$*, where $heat(\xi^i_{\mathrm{u}}, cool) = -1$ if $\xi^i_{\mathrm{u}}(cool) = \mathsf{on}$ (active cooling), and $heat(\xi^i_{\mathrm{u}}, cool) = +1$ if $\xi^i_{\mathrm{u}}(cool) = \mathsf{off}$ (inactive cooling);*
- $\xi_{\mathrm{e}} \in \mathbb{R}^{\{s_{\mathrm{t}}\}}$ *and* $\xi_{\mathrm{e}}(s_{\mathrm{t}}) = 0.1 = \epsilon$*;*
- $meas(\xi^i_{\mathrm{x}}, \xi_{\mathrm{e}}) = \{\xi : \xi(s_{\mathrm{t}}) \in [\xi^i_{\mathrm{x}}(temp) - \epsilon\,, \xi^i_{\mathrm{x}}(temp) + \epsilon]\}$*;*
- $inv(\xi_{\mathrm{x}}) = \mathsf{true}$ *if* $0 \leq \xi_{\mathrm{x}}(temp) \leq 50$*;* $inv(\xi_{\mathrm{x}}) = \mathsf{false}$*, otherwise.*

Let us now formalise the cyber component of CPSs in CCPSA. Our processes build on the *timed process algebra TPL* of [10] (basically CCS enriched with a discrete notion of time). We extend TPL with two main ingredients:

- two different constructs to read values detected at sensors and write values on actuators, respectively;
- special constructs to represent malicious activities on physical devices.

The remaining processes of the calculus are the same as those of TPL.

Definition 3 (Processes)**.** Processes *are defined by the grammar*

$$P, Q \ ::= \ \mathsf{nil} \mid \mathsf{idle}.P \mid P \parallel Q \mid \mathsf{timeout}\lfloor \pi.P \rfloor Q \mid \mathsf{if}\ (b)\ \{P\}\ \mathsf{else}\ \{Q\} \mid P\backslash c \mid X \mid \mathsf{rec}\, X.P.$$

We write nil for the *terminated process*. The process $\mathsf{idle}.P$ sleeps for one time unit and then continues as P. We write $P \parallel Q$ to denote the *parallel composition* of concurrent *threads* P and Q. The process $\mathsf{timeout}\lfloor \pi.P \rfloor Q$, with $\pi \in \{\mathsf{snd}\, c\langle v\rangle, \mathsf{rcv}\, c(x), \mathsf{read}\, s(x), \mathsf{write}\, a\langle v\rangle, \mathsf{read}\, \lightning p(x), \mathsf{write}\, \lightning p\langle v\rangle\}$, denotes *prefixing with timeout*. Thus, $\mathsf{timeout}\lfloor \mathsf{snd}\, c\langle v\rangle.P \rfloor Q$ sends the value v on channel c and, after that, it continues as P; otherwise, after one time unit, it evolves into Q. The process $\mathsf{timeout}\lfloor \mathsf{rcv}\, c(x).P \rfloor Q$ is the obvious counterpart for receiving. $\mathsf{timeout}\lfloor \mathsf{read}\, s(x).P \rfloor Q$ reads the value v detected by the sensor s and, after that, it continues as P, where x is replaced by v; otherwise, after one time unit, it evolves into Q. $\mathsf{timeout}\lfloor \mathsf{write}\, a\langle v\rangle.P \rfloor Q$ writes the value v on the actuator a and, after that, it continues as P; otherwise, after one time unit, it evolves into Q. For $\pi \in \{\mathsf{read}\, \lightning p(x), \mathsf{write}\, \lightning p\langle v\rangle\}$, the process $\mathsf{timeout}\lfloor \pi.P \rfloor Q$ denotes the reading and the writing, respectively, of the physical device p (sensor or actuator) made by the *attack*. Thus, in CCPSA, *attack processes* have specific constructs to interact with physical devices.

The process $P\backslash c$ is the channel restriction operator of CCS. It is quantified over the set $\mathcal{C}$ of communication channels but we sometimes use the shorthand $P\backslash C$ to mean $P\backslash c_1 \backslash c_2 \cdots \backslash c_n$ for $C = \{c_1, c_2, \ldots, c_n\}$.

The process $\mathsf{if}\ (b)\ \{P\}\mathsf{else}\{Q\}$ is the standard conditional, where b is a decidable guard. As in CCS, we assume that

$\mathsf{if}\,(b)\,\{P\}\,\mathsf{else}\,\{Q\} = P$ if $[\![b]\!] = \mathsf{true}$ (i.e., b evaluates to true), and $\mathsf{if}\,(b)\,\{P\}\mathsf{else}\{Q\} = Q$ if $[\![b]\!] = \mathsf{false}$. In processes of the form $\mathsf{idle}.Q$ and $\mathsf{timeout}\lfloor\pi.P\rfloor Q$, the occurrence of Q is said to be *time-guarded*. The process $\mathsf{rec}\,X.P$ denotes *time-guarded recursion* as all occurrences of the process variable X may only occur time-guarded in P.

In the two constructs $\mathsf{timeout}\lfloor\mathsf{rcv}\,c(x).P\rfloor Q$ and $\mathsf{timeout}\lfloor\mathsf{read}\,\mu(x).P\rfloor Q$, with $\mu \in \{p, \not\!p\}$, the variable x is said to be *bound*. Similarly, the process variable X is bound in $\mathsf{rec}\,X.P$. This gives rise to the standard notions of *free/bound (process) variables* and *α-conversion*. A term is *closed* if it does not contain free (process) variables, and we assume to always work with closed processes: the absence of free variables is preserved at run-time. As further notation, we write $T\{^v/_x\}$ for the substitution of the variable x with the value v in any expression T of our language. Similarly, $T\{^P/_X\}$ is the substitution of the process variable X with the process P in T.

The syntax of our CPSs is slightly too permissive as processes and environments might not fit together. We could rule out ill-formed CPSs by means of a simple type system. For simplicity, we prefer to provide the following definition.

Definition 4 (Well-formedness)**.** *Given a process P and an environment $E = \langle \xi_{\mathrm{x}}, \xi_{\mathrm{u}}, \xi_{\mathrm{w}}, evol, \xi_{\mathrm{e}}, meas, inv\rangle$, the CPS $E \bowtie P$ is* well-formed *if: (i) for any sensor s mentioned in P, the function ξ_{e} is defined in s; (ii) for any actuator a mentioned in P, the function ξ_{u} is defined in a; (iii) P doesn't contain prefixes of the form $\mathsf{write}\,s\langle v\rangle$ or $\mathsf{read}\,a(x)$.*

The first two conditions require that the physical devices mentioned in the process must be defined in the environment, whereas the last condition implicitly imposes that only attack processes can write on sensors and read on actuators. Hereafter, we will always work with well-formed networks.

Finally, we assume a number of *notational conventions*. To model *time-persistent prefixing*, we write $\pi.P$ instead of $\mathsf{rec}\,X.\mathsf{timeout}\lfloor\pi.P\rfloor X$, when X does not occur in P. We write $\mathsf{timeout}\lfloor\pi\rfloor Q$ as an abbreviation for $\mathsf{timeout}\lfloor\pi.\mathsf{nil}\rfloor Q$ and $\mathsf{timeout}\lfloor\pi.P\rfloor$ instead of $\mathsf{timeout}\lfloor\pi.P\rfloor\mathsf{nil}$. We write $\mathsf{snd}\,c$ (resp., $\mathsf{rcv}\,c$) when channel c is used for pure synchronisation. For $k \geq 0$, we write $\mathsf{idle}^k.P$ as a shorthand for $\mathsf{idle}.\mathsf{idle}.\ldots\mathsf{idle}.P$, where the prefix idle appears k consecutive times. We write $\mathsf{if}\,(b)\,\{P\}$ instead of $\mathsf{if}\,(b)\,\{P\}\,\mathsf{else}\,\{\mathsf{nil}\}$. Let $M = E \bowtie_{\mathcal{S}} P$, we write $M \parallel Q$ for $E \bowtie_{\mathcal{S}} (P \parallel Q)$.

We can now finalise our running example.

Example 2 (Cyber component of the CPS *Sys*)**.** *Let us define the cyber component of the CPS Sys described in Example 1. We define two parallel processes: Ctrl and IDS. The former models the* controller *activity, consisting in reading the temperature sensor and in governing the cooling system via its actuator, whereas the latter models a simple* intrusion detection system *that attempts to detect and signal abnormal behaviours of the system. Intuitively, Ctrl senses the temperature of the engine at each time slot. When the sensed temperature is above 10 degrees, the controller activates the coolant. The cooling activity is maintained for 5 consecutive time units. After that time, the controller synchronises with the IDS component via a private channel syns, and then waits for* instructions*, via a channel ins. The IDS component checks whether the sensed temperature is still above 10. If this is the case, it sends an* alarm *of "high temperature", via a specific channel, and then says to Ctrl to keep cooling for other 5 time units; otherwise, if the temperature is not above 10, the IDS component requires Ctrl to stop the cooling activity.*

$$
\begin{array}{rcl}
Ctrl & = & \mathsf{rec}\,X.\mathsf{read}\,s_{\mathrm{t}}(x).\mathsf{if}\,(x > 10)\,\{Cooling\}\,\mathsf{else}\,\{\mathsf{idle}.X\} \\
Cooling & = & \mathsf{write}\,\mathsf{cool}\langle\mathsf{on}\rangle.\mathsf{idle}^5.\mathsf{rec}\,Y.\mathsf{snd}\,sync. \\
 & & \mathsf{rcv}\,ins(y).\mathsf{if}\,(y = \mathsf{keep_cooling}) \\
 & & \{\mathsf{idle}^5.Y\}\,\mathsf{else}\,\{\mathsf{write}\,cool\langle\mathsf{off}\rangle.\mathsf{idle}.X\} \\
IDS & = & \mathsf{rec}\,Z.\mathsf{rcv}\,sync.\mathsf{read}\,s_{\mathrm{t}}(x).\mathsf{if}\,(x > 10) \\
 & & \{\mathsf{snd}\,alarm\langle\mathsf{high_temp}\rangle.\mathsf{snd}\,ins\langle\mathsf{keep_cooling}\rangle. \\
 & & \mathsf{idle}.Z\}\,\mathsf{else}\,\{\mathsf{snd}\,ins\langle\mathsf{stop}\rangle.\mathsf{idle}.Z\}
\end{array}
$$

Thus, the whole CPS is defined as:

$$Sys \stackrel{\text{def}}{=} Env \bowtie (Ctrl \parallel IDS)\backslash\{sync, ins\}\,,$$

where Env is the physical environment defined in Example 1. We remark that, for the sake of simplicity, our IDS component is quite basic. However, it is straightforward to replace it with a more sophisticated one, containing more informative tests on sensor values and/or on actuators commands.

We will now define the operational semantics of CCPSA and then we will show that in the absence of attacks, the system *Sys* never sends an *alarm* and never deadlocks (i.e., it never violates the invariant).

2.2. Labelled transition semantics

The physical evolution of a system $E \bowtie_{\mathcal{S}} P$ is given via the evolution map reported in E. We provide the operational semantics of the dynamics of the cyber component P in terms of a *labelled transition system (LTS)* in the SOS style of Plotkin. In Definition 5, for convenience, we define some auxiliary operators on environments.

Definition 5 (Auxiliary operators on environments)**.** *Let $E = \langle \xi_{\mathrm{x}}, \xi_{\mathrm{u}}, \xi_{\mathrm{w}}, evol, \xi_{\mathrm{e}}, meas, inv\rangle$.*

- $read_sensor(E, s) \stackrel{\text{def}}{=} \{\xi(s) : \xi \in meas(\xi_{\mathrm{x}}, \xi_{\mathrm{e}})\}$,
- $update_act(E, a, v) \stackrel{\text{def}}{=} \langle \xi_{\mathrm{x}}, \xi_{\mathrm{u}}[a \mapsto v], \xi_{\mathrm{w}}, evol, \xi_{\mathrm{e}}, meas, inv\rangle$,
- $E\{^{\xi}/_{\xi_{\mathrm{w}}}\} \stackrel{\text{def}}{=} \langle \xi_{\mathrm{x}}, \xi_{\mathrm{u}}, \xi, evol, \xi_{\mathrm{e}}, meas, inv\rangle$,
- $next(E) \stackrel{\text{def}}{=} \bigcup_{\xi \in evol(\xi_{\mathrm{x}}, \xi_{\mathrm{u}}, \xi_{\mathrm{w}})} \langle \xi, \xi_{\mathrm{u}}, \xi_{\mathrm{w}}, evol, \xi_{\mathrm{e}}, meas, inv\rangle$,
- $inv(E) \stackrel{\text{def}}{=} inv(\xi_{\mathrm{x}})$.

The operator $read_sensor(E, s)$ returns the set of possible measurements detected by sensor s in the environment E; it returns a set of possible values rather than a single value due to the error $\xi_{\mathrm{e}}(s)$ of sensor s. $update_act(E, a, v)$ returns the new environment in which the actuator function is updated in such a manner to associate the actuator a with the value v. $E\{^{\xi}/_{\xi_{\mathrm{w}}}\}$ returns a new environment in which the uncertainty function ξ_{w} is replaced by ξ. Given a CPS $M = E \bowtie P$, we often write $M\{^{\xi'_{\mathrm{w}}}/_{\xi_{\mathrm{w}}}\}$ for $E\{^{\xi'_{\mathrm{w}}}/_{\xi_{\mathrm{w}}}\} \bowtie P$. $next(E)$ returns the set of the next admissible environments

$$\text{(Outp)}\ \frac{-}{\mathsf{timeout}\lfloor \mathsf{snd}\, c\langle v\rangle.P\rfloor Q \xrightarrow{\overline{c}v} P} \qquad \text{(Inpp)}\ \frac{-}{\mathsf{timeout}\lfloor \mathsf{rcv}\, c(x).P\rfloor Q \xrightarrow{cv} P\{^{v}/_{x}\}} \qquad \text{(Com)}\ \frac{P \xrightarrow{\overline{c}v} P' \quad Q \xrightarrow{cv} Q'}{P \parallel Q \xrightarrow{\tau} P' \parallel Q'}$$

$$\text{(Write)}\ \frac{\mu \in \{p, \lightning p\}}{\mathsf{timeout}\lfloor \mathsf{write}\, \mu\langle v\rangle.P\rfloor Q \xrightarrow{\mu!v} P} \qquad \text{(Read)}\ \frac{\mu \in \{p, \lightning p\}}{\mathsf{timeout}\lfloor \mathsf{read}\, \mu(x).P\rfloor Q \xrightarrow{\mu?v} P\{^{v}/_{x}\}}$$

$$\text{(}\lightning\text{SensWrite}\lightning\text{)}\ \frac{P \xrightarrow{\lightning s!v} P' \quad Q \xrightarrow{s?v} Q'}{P \parallel Q \xrightarrow{\tau:s} P' \parallel Q'} \qquad \text{(}\lightning\text{ActRead}\lightning\text{)}\ \frac{P \xrightarrow{a!v} P' \quad Q \xrightarrow{\lightning a?v} Q'}{P \parallel Q \xrightarrow{\tau:a} P' \parallel Q'} \qquad \text{(Par)}\ \frac{P \xrightarrow{\lambda} P' \quad \lambda \neq \mathsf{idle}}{P \parallel Q \xrightarrow{\lambda} P' \parallel Q}$$

$$\text{(Res)}\ \frac{P \xrightarrow{\lambda} P' \quad \lambda \notin \{cv, \overline{c}v\}}{P\backslash c \xrightarrow{\lambda} P'\backslash c} \qquad \text{(Rec)}\ \frac{P\{^{\mathsf{rec}\, X.P}/_{X}\} \xrightarrow{\lambda} Q}{\mathsf{rec}\, X.P \xrightarrow{\lambda} Q} \qquad \text{(TimeNil)}\ \frac{-}{\mathsf{nil} \xrightarrow{\mathsf{idle}} \mathsf{nil}}$$

$$\text{(Delay)}\ \frac{-}{\mathsf{idle}.P \xrightarrow{\mathsf{idle}} P} \qquad \text{(Timeout)}\ \frac{-}{\mathsf{timeout}\lfloor \pi.P\rfloor Q \xrightarrow{\mathsf{idle}} Q} \qquad \text{(TimePar)}\ \frac{P \xrightarrow{\mathsf{idle}} P' \quad Q \xrightarrow{\mathsf{idle}} Q'}{P \parallel Q \xrightarrow{\mathsf{idle}} P' \parallel Q'}$$

TABLE 1. LTS FOR PROCESSES

reachable from E, by an application of the evolution map. $inv(E)$ checks whether the invariant is satisfied by the current values of the state variables (here, with an abuse of notation, we overload the meaning of the function inv).

In Table 1, we provide transition rules for processes. Here, the meta-variable λ ranges over labels in the set $\{\mathsf{idle}, \tau, \overline{c}v, cv, a!v, s?v, \lightning p!v, \lightning p?v, \tau{:}p\}$. Rules (Outp), (Inpp) and (Com) serve to model channel communication, on some channel c. Rules (Write) and (Read) denote the writing/reading of some data on the physical device p. Rule (⚡SensWrite⚡) models an *integrity attack on sensor* s, where the controller of s is supplied with a fake value v provided by the attack. Rule (⚡ActRead⚡) models a *DoS attack to the actuator* a, where the update request of the controller is intercepted by the attacker and it never reaches the actuator. Rule (Par) propagates untimed actions over parallel components. Rules (Res) and (Rec) are standard rules for channel restriction and recursion, respectively. The following four rules model the passage of one time unit. The symmetric counterparts of rules (Com) and (Par) are obvious and thus omitted from the table.

In Table 2, we lift the transition rules from processes to systems. Except for the last one, all rules have a common premise $inv(E)$: a system can evolve only if the invariant is satisfied. Here, actions, ranged over by α, are in the set $\{\tau, \overline{c}v, cv, \mathsf{idle}, \mathsf{deadlock}\}$. These actions denote: non-observable activities (τ); observable logical activities, more precisely, channel transmission ($\overline{c}v$ and cv); the passage of time (idle); and a specific observable physical event, system deadlock (deadlock). Rules (Out) and (Inp) model transmission and reception, with an external system, on a channel c. Rule (SensReadSec) models the reading of the current data detected at a *secured sensor* s, whereas rule (SensReadUnsec) models the reading of an *unsecure sensor* s. In this case, since the sensor is not protected, the presence of a malicious action $\lightning s!w$ prevents the reading of the sensor. We already said that rule (⚡SensWrite⚡) of Table 1 models integrity attacks on an unsecure sensor s, however, together with rule (SensReadUnsec), it also serves to model *DoS attacks on an unsecure sensor* s, as the controller of s cannot read its correct value if the attacker is currently supplying a fake value for it.

Rule (⚡SensRead⚡) allows the attacker to read the confidential value detected at an unsecure sensor s. Rule (ActWriteSec) models the writing of a value v on a secured actuator a, whereas rule (ActWriteUnsec) models the writing on a unsecured actuator a. Again, if the actuator is unprotected, the presence of an attack (capable of performing an action $\lightning a?v$) prevents the correct access to the actuator by the controller. Rule (⚡ActWrite⚡) models an *integrity attack to an unsecure actuator* a, where the attack updates the actuator with a fake value.

Rule (Tau) lifts non-observable actions from processes to systems. This includes communications channels and attacks' accesses to unsecured physical devices. A similar lifting occurs in rule (Time) for timed actions, where $\mathrm{next}(E)$ returns the set of possible environments for the next time slot. Thus, by an application of rule (Time) a CPS moves to the next physical state, in the next time slot. Finally, rule (Deadlock) is introduced for convenience to directly observe the violation of the invariant. We know that when the invariant is violated, a system deadlock occurs. In CCPSA, when the invariant is violated, the system emits a special action deadlock, forever.

Now, having defined the actions that can be performed by a system, we can easily concatenate these actions to define the possible execution traces of the system. Formally, given a trace of actions $t = \alpha_1 \dots \alpha_n$, we will write $\xrightarrow{t}$ as an abbreviation for $\xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n}$. In the rest of the paper, we will use the function $\#\mathsf{idle}(t)$ to get the number of occurrences of the timed action idle in the trace t.

Let us apply our notion of trace to formalise properties of our example system Sys. The following proposition shows that Sys never deadlocks and never transmits on the channel $alarm$.

Proposition 1. *Let Sys be the CPS defined in Example 2. If $Sys \xrightarrow{t} Sys'$, for some trace $t = \alpha_1 \dots \alpha_n$, then $\alpha_i \in$*

$$\text{(Out)}\ \frac{P \xrightarrow{\overline{c}v} P' \quad \mathrm{inv}(E)}{E \bowtie_{\mathcal{S}} P \xrightarrow{\overline{c}v} E \bowtie_{\mathcal{S}} P'} \qquad \text{(Inp)}\ \frac{P \xrightarrow{cv} P' \quad \mathrm{inv}(E)}{E \bowtie_{\mathcal{S}} P \xrightarrow{cv} E \bowtie_{\mathcal{S}} P'}$$

$$\text{(SensReadSec)}\ \frac{P \xrightarrow{s?v} P' \quad s \in \mathcal{S} \quad \mathrm{inv}(E) \quad v \in \mathrm{read_sensor}(E,s)}{E \bowtie_{\mathcal{S}} P \xrightarrow{\tau} E \bowtie_{\mathcal{S}} P'}$$

$$\text{(SensReadUnsec)}\ \frac{P \xrightarrow{s?v} P' \quad s \notin \mathcal{S} \quad P \overset{\lightning_{s!v}}{\not\longrightarrow} \quad \mathrm{inv}(E) \quad v \in \mathrm{read_sensor}(E,s)}{E \bowtie_{\mathcal{S}} P \xrightarrow{\tau} E \bowtie_{\mathcal{S}} P'}$$

$$\text{(}\lightning\text{SensRead}\lightning\text{)}\ \frac{P \xrightarrow{\lightning_{s?v}} P' \quad s \notin \mathcal{S} \quad \mathrm{inv}(E) \quad v \in \mathrm{read_sensor}(E,s)}{E \bowtie_{\mathcal{S}} P \xrightarrow{\tau} E \bowtie_{\mathcal{S}} P'}$$

$$\text{(ActWriteSec)}\ \frac{P \xrightarrow{a!v} P' \quad a \in \mathcal{S} \quad \mathrm{inv}(E) \quad E' = \mathrm{update_act}(E,a,v)}{E \bowtie_{\mathcal{S}} P \xrightarrow{\tau} E' \bowtie_{\mathcal{S}} P'}$$

$$\text{(ActWriteUnsec)}\ \frac{P \xrightarrow{a!v} P' \quad a \notin \mathcal{S} \quad P \overset{\lightning_{a?v}}{\not\longrightarrow} \quad \mathrm{inv}(E) \quad E' = \mathrm{update_act}(E,a,v)}{E \bowtie_{\mathcal{S}} P \xrightarrow{\tau} E' \bowtie_{\mathcal{S}} P'}$$

$$\text{(}\lightning\text{ActWrite}\lightning\text{)}\ \frac{P \xrightarrow{\lightning_{a!v}} P' \quad a \notin \mathcal{S} \quad \mathrm{inv}(E) \quad E' = \mathrm{update_act}(E,a,v)}{E \bowtie_{\mathcal{S}} P \xrightarrow{\tau} E' \bowtie_{\mathcal{S}} P'}$$

$$\text{(Tau)}\ \frac{(P \xrightarrow{\tau} P') \vee (P \xrightarrow{\tau:p} P' \ p \notin \mathcal{S}) \quad \mathrm{inv}(E)}{E \bowtie_{\mathcal{S}} P \xrightarrow{\tau} E \bowtie_{\mathcal{S}} P'} \qquad \text{(Deadlock)}\ \frac{\neg\, \mathrm{inv}(E)}{E \bowtie_{\mathcal{S}} P \xrightarrow{\mathsf{deadlock}} E \bowtie_{\mathcal{S}} P}$$

$$\text{(Time)}\ \frac{P \xrightarrow{\mathsf{idle}} P' \quad E \bowtie_{\mathcal{S}} P \overset{\tau}{\not\longrightarrow} \quad \mathrm{inv}(E) \quad E' \in \mathrm{next}(E)}{E \bowtie_{\mathcal{S}} P \xrightarrow{\mathsf{idle}} E' \bowtie_{\mathcal{S}} P'}$$

TABLE 2. LTS FOR CPSS

$\{\tau, \mathsf{idle}\}$, *for any* $i \in \{1, \ldots, n\}$.

Actually, we can be quite precise on the temperature reached by *Sys* before and after the cooling activity: in each of the 5 rounds of cooling, the temperature will drop of a value laying in the interval $[1-\delta, 1+\delta]$, where δ is the uncertainty of the model.

Proposition 2. *Let Sys be the CPS defined in Example 2. For any execution trace of Sys, we have:*

- *when Sys* turns on *the cooling, the value of the state variable temp ranges over* $(9.9\,,\, 11.1 + \delta]$*;*
- *when Sys* turns off *the cooling, the value of the variable temp ranges over* $(9.9 - 5*(1+\delta)\,,\, 11.1 + \delta - 5*(1-\delta)]$.

In Figure 2, the left graphic collects a campaign of 100 simulations, lasting 250 time units each, showing that the value of the state variable *temp* when the cooling system is turned on (respectively, off) lays in the interval $(9.9, 11.5]$ (respectively, $(2.9, 8.5]$); these bounds are represented by the dashed horizontal lines. Since $\delta = 0.4$, these results are in line with those of Proposition 2. The right graphic shows three examples of possible evolutions in time of the state variable *temp*.

2.3. Behavioural semantics

Having defined a labelled transition semantics, we are ready to formalise a notion of observational semantics for CPSs, based on execution traces.

We adopt a standard notation for weak transitions: we denote with $\Rightarrow$ the reflexive and transitive closure of τ-actions, namely $(\xrightarrow{\tau})^*$, whereas $\overset{\alpha}{\Longrightarrow}$ means $\Longrightarrow \xrightarrow{\alpha} \Longrightarrow$, and finally $\overset{\hat{\alpha}}{\Longrightarrow}$ denotes $\Rightarrow$ if $\alpha = \tau$ and $\overset{\alpha}{\Longrightarrow}$ otherwise. Given a trace of actions $t = \alpha_1 \ldots \alpha_n$, write $\overset{\hat{t}}{\Longrightarrow}$ as an abbreviation for $\overset{\hat{\alpha_1}}{\Longrightarrow} \ldots \overset{\hat{\alpha_n}}{\Longrightarrow}$.

Definition 6 (Trace equivalence)**.** *Given CPSs M and N, we write $M \sqsubseteq N$ if $M \xrightarrow{t} M'$ implies that there exists N' such that $N \overset{\hat{t}}{\Longrightarrow} N'$, for any trace t. If $M \sqsubseteq N$ and $N \sqsubseteq M$, then we say that M and N are* trace equivalent *and write $M \simeq N$.*

Remark 1. *Unlike standard trace semantics, due to the presence of the rule (Deadlock) and the special action* deadlock*, our trace equivalence is able distinguish a CPS that may* deadlock *from one that may not.*

We recall that the only *observable activities* in CCPSA are: time passing, channel communication and system deadlock. As a consequence, the capability to observe physical events different from deadlocks depends of the capability of

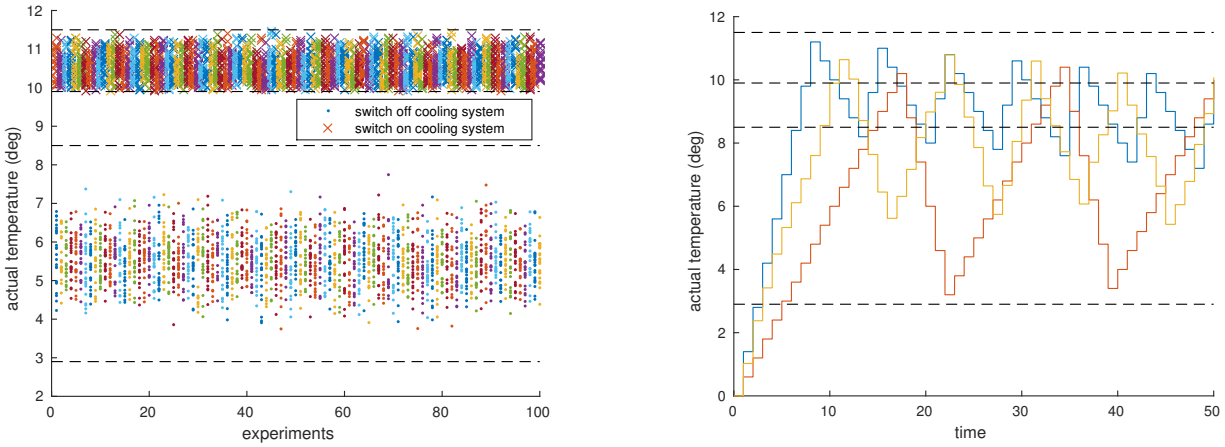

Figure 2. Simulations on the CPS of Example 2

the cyber component to recognise those events by acting on sensors and actuators, and then signalling those events via communications on (unrestricted) channels.

As we are interested in examining timing aspects of attacks, such as beginning and duration, we propose a variant of trace equivalence between two CPSs up to (a possibly infinite) time interval.

Definition 7 (Trace equivalence up to a time interval)**.** *Let M and N be CPSs, $m \in \mathbb{N}^+$ and $n \in \mathbb{N}^+ \cup \infty$. We write $M \sqsubseteq_{[m,n]} N$ if the following conditions hold:*

- *m is the minimum integer for which there is a trace t, with $\#\mathsf{idle}(t) = m-1$, such that $M \xrightarrow{t}$ and $N \stackrel{\hat{t}}{\Longrightarrow}$;*
- *n is the minimum element of $\mathbb{N}^+ \cup \infty$, $n \geq m$, such that whenever $M \xrightarrow{t} M'$, with $\#\mathsf{idle}(t) = n-1$, there is t', with $\#\mathsf{idle}(t') = \#\mathsf{idle}(t)$, such that $N \xrightarrow{t'} N'$ and $M' \sqsubseteq N'$.*

We write $M \simeq_{[m,n]} N$ if $M \sqsubseteq_{[m_1,n_1]} N$, $N \sqsubseteq_{[m_2,n_2]} M$ with $m = \min(m_1, m_2)$ and $n = \max(n_1, n_2)$.

Intuitively, if $M \sqsubseteq_{[m,n]} N$, then N simulates the execution traces of M, except for the time interval $[m, n]$. In the rest of the paper, we assume that $\sqsubseteq_{[\infty,\infty]} \stackrel{\text{def}}{=} \sqsubseteq$ and $\simeq_{[\infty,\infty]} \stackrel{\text{def}}{=} \simeq$.

Finally, note that we could have equipped CCPSA with a *(bi)simulation-based* behavioural semantics rather than a trace-based one; however, since our trace semantics is sensitive to deadlocks of the physical plant, it proves to be fully adequate for the purposes of this paper.

3. Cyber-Physical Attacks

In this section, we use CCPSA to formalise a *threat model* where attacks can manipulate sensor and/or actuator signals in order to drive a CPS into an undesired state [24]. An attack may have different levels of access to physical devices depending on the model assumed. For example, it might be able to get read access to the sensors but not write access; or it might get write-only access to the actuators but not read-access. As observed in [4], this level of granularity is very important to model precisely how attacks can affect a CPS. For simplicity, in this paper we don't represent attacks on communication channels as our focus is on attacks to physical devices.

The syntax of our cyber-physical attack is a slight restriction of that seen for processes: in terms of the form $\mathsf{timeout}\lfloor \pi.P \rfloor Q$, we require $\pi \in \{\mathsf{write}\,\lightning p\langle v\rangle, \mathsf{read}\,\lightning p(x)\}$. In this manner, we provide a syntactic way to distinguish attacks from genuine processes.

Definition 8 (Honest system)**.** *A CPS $E \bowtie_{\mathcal{S}} P$ is* honest *if P is honest, where a process P is honest if it does not contain prefixes of the form* $\mathsf{write}\,\lightning p\langle v\rangle$ *or* $\mathsf{read}\,\lightning p(x)$.

The following proposition formalises the preemptive power ensured by our operational semantics to prefixes of the form $\mathsf{write}\,\lightning p\langle v\rangle$ and $\mathsf{read}\,\lightning p(x)$ *on unsecure devices* p: an attack process in CCPSA can always prevent the regular access to a physical device (sensor or actuator) by its controller.

Proposition 3 (Attack preemptiveness)**.** *Let $M = E \bowtie_{\mathcal{S}} P$ be a CPS.*

- *If there is Q such that $P \xrightarrow{\lightning s!v} Q$, with $s \notin \mathcal{S}$, then there is no M' such that $M \xrightarrow{\tau} M'$ by an application of the rule (SensReadUnsec).*
- *If there is Q such that $P \xrightarrow{\lightning a?v} Q$, with $a \notin \mathcal{S}$, then there is no M' such that $M \xrightarrow{\tau} M'$ by an application of the rule (ActWriteUnsec).*

We group attacks in classes that describe both the malicious activity and the timing aspects of the attack. Thus, let $\mathcal{I}$ be a set of malicious activities on a number of physical devices, $m \in \mathbb{N}^+$ be the time slot when an attack starts, and $n \in \mathbb{N}^+ \cup \infty$ be the time slot when the attack ends; we say that an *attack A is of class $\mathcal{I}[m, n]$* if: (1) all possible malicious actions of A coincide with those contained in $\mathcal{I}$, (2) the first of those actions may occur in the m-th time slot

(i.e., after $m-1$ idle-actions), and (3) the last of those actions may occur in the n-th time slot (i.e., after $n-1$ idle-actions). Formally:[2]

Definition 9 (Class of attacks)**.** *Let $\mathcal{I} \subseteq \{\lightning p? : p \in \mathcal{S} \cup \mathcal{A}\} \cup \{\lightning p! : p \in \mathcal{S} \cup \mathcal{A}\}$ be a set of* malicious activities*, $m \in \mathbb{N}^+$, and $n \in \mathbb{N}^+ \cup \infty$. An attack A is of* class $\mathcal{I}[m,n]$ *whenever:*

1) *$\lightning p? \in \mathcal{I}$ (respectively, $\lightning p! \in \mathcal{I}$) iff there is a trace t containing $\lightning p?v$, for some v, (respectively, $\lightning p!v$, for some v) such that $A \xrightarrow{t} A'$;*
2) *$m-1 = \inf\{ \#\mathsf{idle}(t) : A \xrightarrow{t} \xrightarrow{\iota v} A' \text{ with } \iota \in \mathcal{I} \}$;*
3) *$n-1 = \sup\{ \#\mathsf{idle}(t) : A \xrightarrow{t} \xrightarrow{\iota v} A' \text{ with } \iota \in \mathcal{I} \}$.*

According to the approach proposed in [7], we can say that an attack A affects a CPS M if the execution of the composed system $M \parallel A$ differs from that of the original system M, in an observable manner. Basically, the attack can influence the system under attack in three different ways:

- The system $M \parallel A$ might deadlock when M may not; this means that the attack A affects the *availability* of the system. We recall that in the context of CPSs, deadlock is a particular sever physical event.
- The system $M \parallel A$ might have non-genuine execution traces containing observables that cannot be reproduced by M; here the attack affects the *integrity* of the system behaviour.
- The system M might have execution traces containing observables that cannot be reproduced by the system under attack $M \parallel A$ (because they are prevented by the attack); again this is an attack against the *availability* of the system (and to some extent also to its integrity).

Definition 10 (Attack tolerance/vulnerability)**.** *An honest CPS M is* tolerant to an attack *A if and only if $M \parallel A \simeq M$. An honest CPS M is* vulnerable to an attack *A if and only if there is a time interval $[m,n]$, with $m \in \mathbb{N}^+$ and $n \in \mathbb{N}^+ \cup \infty$, such that $M \parallel A \simeq_{[m,n]} M$.*

According to Definition 10, the tolerance/vulnerability of a CPS to an attack strongly depends on the capability of its cyber component to recognise abnormal physical states. Note also that if a system M is vulnerable to an attack A of class $\mathcal{I}[m,n]$ in the time interval $[m',n']$ (where, obviously, $m' \geq m$) then: (i) n is finite iff the attack A *terminates*; (ii) the attack yields observable effects in the time interval $[m',n']$. Thus, if n' is finite we have a *temporary attack*, otherwise we have a *permanent attack*. If $m' > n$, then we have a *stealthy attack*, i.e., an attack that produces observable effects only after its termination [24]. Moreover, if $m' - n$ is big enough and $n - m$ is small, then we have a quick nasty attack that affects the system late enough to allow *attack camouflages* [9]. On the other hand, if m' is significantly smaller than n, then the attack affects the observable behaviour of the system well before its termination and the CPS has good chances of undertaking countermeasures to stop the attack.

2. For simplicity, Definition 9 relies on the operational semantics of the language. A more effective way to define classes of attacks could make use of a simple type system.

In the rest of this section we present a number of different attacks to the CPS Sys described in Example 2.

Example 3. *Consider the following* DoS/Integrity attack*, of class $\{\lightning cool?, \lightning cool!\}[m,m]$, on the controller of the actuator cool:*

$$A_m \;=\; \mathsf{idle}^{m-1}.\mathsf{timeout}\lfloor \mathsf{read}\,\lightning cool(x).\mathsf{if}\ (x = \mathsf{off})\ \{\mathsf{write}\,\lightning cool\langle \mathsf{off}\rangle\}\rfloor\ .$$

Here, the attack operates exclusively in the m-th time slot, when she tries to steal the cooling command (on or off) coming from the controller, and fabricates a fake command to turn off the cooling system. In practice, if the controller sends a command to turn off the coolant, nothing bad will happen as the attack will put the same message back. If the controller sends in the m-th time slot a command to turn the cooling on, then this command will be eavesdropped by the attack triggering an alarm after 5 time units. Note that the command to turn the cooling on is never re-sent by Ctrl. As a consequence, the temperature will continue to rise and the IDS component will keep sending alarms every 5 time units, until the whole system deadlocks because the temperature reaches the threshold of 50 degrees. Note also that, since the initial temperature of the system is 0, the cooling system may not be activated before the 9-th time slot (i.e., before 8 idle-actions). Thus, for $m \leq 8$ the attack is harmless, whereas for $m > 8$ we have to wait for the termination of the cooling cycle to possibly observe the first alarm fired by the IDS. Thus, this is an example of stealthy attack. Finally, note that if the controller does not send any command to the cooling system in the m-th time slot, then the attack is harmless.

Proposition 4. *Let Sys be the CPS defined in Example 2, and A_m be the attack defined in Example 3. Then:*

- *$Sys \parallel A_m \simeq Sys$, for $m \leq 8$,*
- *$Sys \parallel A_m \simeq_{[m+5,\infty]} Sys$, for $m > 8$.*

Note that if we considered a simpler DoS attack:

$$\mathsf{idle}^{m-1}.\mathsf{timeout}\lfloor \mathsf{read}\,\lightning cool(x)\rfloor\,,$$

then (for $m \geq 13$) this attack might cause a deadlock of the system if the stolen message were a command to turn off the cooling. In that case, the temperature would continue to drop under the effect of the coolant, until the invariant of the environment is violated. Obviously, a more accurate IDS component could detect the dropping of the temperature before reaching the deadlock of the whole system.

Example 4. *Consider the following* DoS/Integrity attack*, of class $\{\lightning s_{\mathrm{t}}!, \lightning s_{\mathrm{t}}?\}[2,\infty]$, to the controller of sensor s_{t}:*

$$A \;=\; \mathsf{idle}.\mathsf{timeout}\lfloor \mathsf{read}\,\lightning s_{\mathrm{t}}(x).\ \mathsf{rec}\,X.\mathsf{timeout}\lfloor \mathsf{write}\,\lightning s_{\mathrm{t}}\langle x\rangle.\mathsf{idle}.X \rfloor X \rfloor\,.$$

Here, the attack A does the following actions in sequence: (i) she sleeps for one time unit, (ii) in the following time slot, she reads the current temperature v at sensor s_{t}, and (iii) for the rest of her life, she keeps sending the same temperature v to the controller of s_{t}.

In the presence of this attack, the process Ctrl never activates the Cooling component (and, hence, nor the IDS component, which is the only one which could send an alarm) as it will always detect a temperature below 10*. Thus, the system Sys* $\parallel$ *A will eventually move to a state in which the invariant is violated and the system will deadlock. This is because there is an execution trace in which after only* $\lceil \frac{50}{1+\delta} \rceil = \lceil \frac{50}{1.4} \rceil = 36$ *time units, in the* 37*-th time slot, the invariant will be broken because the state variable temp reaches* 50.4 *degrees. Note that this is the quickest trace breaking the upper bound of the invariant: the CPS cannot reach a temperature above* 50 *before the* 37*-th time slot. Thus, Sys* $\parallel A \simeq_{[37,\infty]} Sys$*.*

The high number of battery-powered physical devices, possibly interconnected through wireless networks, prevents the encryption/decryption of all packets due to energy constraints. So, the CPS designer may get to a point where she has to understand whether a device is *critical* to the correct functioning of the whole CPS and hence deserves to be protected. In other words, the designer will have to decide which of the many devices to secure using cryptography.

Our semantics ensures that secured devices cannot be attacked, as stated by the following proposition.

Proposition 5. *Let* $M = E \bowtie_{\mathcal{S}} P$ *be an honest CPS. Then, for any class of attacks* $\mathcal{I}[m,n]$*, such that* $\{p : \lightning p? \in \mathcal{I} \text{ or } \lightning p! \in \mathcal{I}\} \subseteq \mathcal{S}$*, we have that M tolerates all attacks A of class* $\mathcal{I}[m,n]$*. Formally,* $M \parallel A \simeq M$*, for any of such attacks A.*

In the following example, we compose the two previous attacks to make a few considerations when more devices are under attack. In particular, the example also shows that deciding whether or not it is worth to secure a device may be a crucial design choice.

Example 5. *Suppose that the designer of Sys knows that the system is exposed to the* DoS/Integrity attack

$$B \;=\; A_{100} \parallel A\,,$$

where A_{100} *is the attack of Example 3, for* $m = 100$*, and A is the attack of Example 4.*

As the attack A_{100} *sleeps for the first* 99 *time slots, the attack A prevails over* A_{100}*, and the reasoning carried out in Example 4 can be adapted to the attack B to derive that Sys* $\parallel B \simeq_{[37,\infty]} Sys$*. This means that the attack B produces observable effects (in this case a deadlock of the whole system) in the* 37*-th time slot.*

Here, an eventual protection of the device cool would not help, because

$$Sys' \parallel B \simeq_{[37,\infty]} Sys'$$

when $Sys' = Env \bowtie_{\{cool\}} Ctrl$*. On the other hand, protecting the sensor* s_{t} *does make a difference. At least in terms of timing. This is because*

$$Sys'' \parallel B \simeq_{[105,\infty]} Sys''$$

for $Sys'' = Env \bowtie_{\{s_{\mathrm{t}}\}} Ctrl$*. In fact, by Proposition 5 the attack A on sensor* s_{t} *is harmless on Sys″. Moreover, by adapting the proof of Proposition 4, the attack* A_{100} *can eavesdrop the cooling command of Sys″ in the* 100*-th time slot. This will lead to an alarm after* 5 *time slots, i.e., in the* 105*-th time slot. In this situation, as we already said in Example 3, the IDS component will keep sending alarms every* 5 *time units, until the whole system deadlocks.*

In a similar case, the CPS designer will probably choose to protect the sensor s_{t} *exposing the CPS to an attack that gives rise to a number of alarms before the whole system gets compromised.*[3]

Let us examine a new attack.

Example 6. *Consider the following DoS/Integrity attack, of class* $\{\lightning s_{\mathrm{t}}!, \lightning s_{\mathrm{t}}?\}[1,n]$*, to the controller of sensor* s_{t}*:*

$$\begin{array}{lcl} A_0 & = & \mathsf{nil} \\ A_n & = & \mathsf{timeout} \lfloor \mathsf{read}\, \lightning s_{\mathrm{t}}(x).\mathsf{timeout} \lfloor \mathsf{write}\, \lightning s_{\mathrm{t}} \langle x{-}2 \rangle . \\ & & \mathsf{idle}.A_{n-1} \rfloor A_{n-1} \rfloor A_{n-1} \end{array}$$

In this attack, for n *consecutive time slots,* A_n *sends to the controller the current sensed temperature decreased by the constant* 2*. The impact of this attack on the system depends on the duration* n *of the attack itself. In particular,*

- *for* $n \leq 8$*, the attack is harmless as the system may not reach a critical temperature above* 10 *in the first* 8 *time slots;*
- *for* $n = 9$*, the system might reach a temperature above* 10 *in the* 9*-th time slot, and the attack would delay the activation of the cooling system of one time slot; however this delay will be absorbed by the system without any observable consequence;*
- *for* $10 \leq n \leq 15$*, the attack is tolerated for at most* 15 *time slots; then, in one of the time slots belonging to the interval* $[16, n{+}6]$ *the system under attack may send an alarm.*
- *for* $n \geq 16$*, the attack is tolerated for at most* n *time slots; then, in one of the time slots belonging to the interval* $[n{+}1, n{+}6]$ *the system under attack may send an alarm.*

Note that, for $n \geq 10$*,* A_n *is an example of* stealthy attack*. This is also a* transitory attack *as the cooling system will safely drop the temperature below* 10 *degrees in at most* 5 *time units.*

Proposition 6. *Let Sys be the CPS defined in Example 2, and* A_n *be the attack defined in Example 6. Then:*

- $Sys \parallel A_n \simeq Sys$*, for* $n \leq 9$*,*
- $Sys \parallel A_n \simeq_{[16,n+6]} Sys$*, for* $10 \leq n \leq 15$*,*
- $Sys \parallel A_n \simeq_{[n+1,n+6]} Sys$*, for* $n \geq 16$*.*

At this point, the reader may wonder whether the designer of the logic of a CPS may take some countermeasures to prevent those kinds of attack by providing a better and more robust logic for their system. More precisely, given a CPS

3. Obviously, our cyber component can be made more robust by re-sending, for instance, the command to turn on the cooling each time the alarm is sent. However, as shown Example 7, in general, similar changes might make the system more robust with respect to an attack but weaker with respect to another.

$M = E \bowtie P$ that is subject to an attack A, is it possible to find a process P' such that the CPS $M' = E \bowtie P'$ is trace equivalent to M but it is not subject to the attack A? In some cases, the answer to this question is "yes"!

Example 7. *Let us consider a variant of the process $Ctrl$ of Example 2:*

$$Ctrl' \;=\; \mathsf{rec}\, X.\mathsf{read}\, s_{\mathrm{t}}(x).\mathsf{if}\,(x > 9)\,\{Cooling\}\,\mathsf{else}\,\{\mathsf{idle}.X\}\,.$$

The process $Ctrl'$ anticipates the activation of the coolant when the sensed temperature is greater than 9. However, the IDS component remains unchanged as it sends an alarm only if the sensed temperature is above 10 degrees after 5 time units of cooling.

Let Sys' be the CPS obtained by replacing $Ctrl$ with $Ctrl'$. Unlike Sys, system Sys' tolerates the attack A_n of class $\{\lightning s_{\mathrm{t}}!, \lightning s_{\mathrm{t}}?\}[1,n]$ of Example 6. Formally, $Sys' \parallel A_n \simeq Sys'$, for any $n \in \mathbb{N}$. This result relies on three facts:

- *Sys and Sys' are trace equivalent. It is easy to show that Proposition 1 holds when replacing Sys with the more prudent version Sys'.*
- *$Sys' \parallel A_n$ may never emit an alarm. In fact, since $\delta = 0.4$ and $\epsilon = 0.1$, when the sensed temperature is greater than 9, by a slight adaptation of Proposition 2 it follows that the value of the state variable $temp$ is in the interval $(10.9, 12.5]$. Thus, after 5 time units of cooling, the value of the state variable $temp$ will be at most $12.5 - 5 * (1 - \delta) = 12.5 - 5 * (.6) = 9.5$, and hence no alarm will fire (recall that Sys', like Sys, sends an alarm only if the sensed temperature is above 10).*
- *$Sys' \parallel A_n$ may never deadlock. As the CPS may never emit an alarm, it follows that the cooling system is always successful in dropping the temperature below 10. Thus, the temperature may never reach 50 degrees. Furthermore, since the attack does not directly operate on the actuator $cool$, after 5 time units of cooling the cooling system is always turned off. In that moment, the temperature is at least $9 - 5*(1 + \delta) = 2$ degrees. As a consequence, the temperature may never get below 0, and thus, since the invariant is never violated, $Sys' \parallel A_n$ may never deadlock.*

So, we may be induced to believe that Sys' is somehow more secure than Sys. However, things are not that simple, and designing a secure variant of a CPS is a complex task that calls for *formal definitions* and *automatic support*. In fact, a different implementation of the logic might be more robust with respect to a certain attack but it could open the way to new attacks, even of the same class.

For instance, Sys' is not more secure than Sys with respect to attacks of class $\{\lightning s_{\mathrm{t}}!, \lightning s_{\mathrm{t}}?\}[1,n]$. As an example, consider the following *Integrity attack* A of class $\{\lightning s_{\mathrm{t}}!, \lightning s_{\mathrm{t}}?\}[1,16]$:

$$\begin{array}{rcl} A & = & \mathsf{timeout}\lfloor \mathsf{read}\, \lightning s_{\mathrm{t}}(x).\mathsf{idle}.A_1 \rfloor \\ A_i & = & \mathsf{timeout}\lfloor \mathsf{read}\, \lightning s_{\mathrm{t}}(x).\mathsf{idle}.A_{i+1} \rfloor, \text{ for } 1 \le i \le 9 \\ A_{10} & = & \mathsf{timeout}\lfloor \mathsf{read}\, \lightning s_{\mathrm{t}}(x). \\ & & \quad \mathsf{if}\,(x = 9.1)\,\{\mathsf{idle}^5.\mathsf{timeout}\lfloor \mathsf{write}\, \lightning s_{\mathrm{t}}\langle 10.5\rangle \rfloor\}\rfloor\,. \end{array}$$

Here, in the first 10 time slots the attack does an harmless reading of the sensor s_{t}, then, in the 11-th time slot, she reads the temperature and checks whether it is equal to 9.1. If this is the case, she waits for 5 time units (when the cooling might be active) and then, in the 16-th time slot, she sends a fake value to the controller of the sensor. This attack is tolerated by Sys but not by Sys'. This is a typical example of an attack that is successful only because it operates when the target system Sys' is very close to a sensitive threshold [14].

Proposition 7. *Let Sys be the CPS defined in Example 2, Sys'' be its variant defined in Example 7, and A be the attack defined in Example 7. Then:*

- $Sys \parallel A \;\simeq\; Sys$,
- $Sys' \parallel A \;\simeq_{[16,\infty]}\; Sys'$.

4. Impact of an attack

In the previous section, we have grouped cyber-physical attacks by focussing on the physical devices under attack and the timing aspects of the attack (Definition 9). Then, we have provided a formalisation of when a CPS should be considered tolerant/vulnerable to an attack (Definition 10).[4] In this section, we will show that it is important not only to demonstrate the tolerance (or vulnerability) of a CPS with respect to certain attacks, but also to evaluate the disruptive impact of those attacks on the CPS [8]. This evaluation is of great importance in big CPSs. For instance, many industrial systems continue their activities even under attack, as the cost of stopping the system would be much higher than the damage caused by the attack itself.

The goal of this section is twofold: to provide a metric to estimate the impact of a successful attack on a CPS, and to investigate possible quantifications of the chances for an attack of being successful when attacking a CPS.

For what concerns the metric, we focus on the ability that an attack may have to drag a CPS out of the correct behaviour modelled by its evolution map, with the given uncertainty. We recall that in Definition 2 the evolution map of our CPSs is naturally monotone with respect to the uncertainty, i.e., if $\xi_{\mathrm{w}} \le \xi'_{\mathrm{w}}$, then $evol(\xi_{\mathrm{x}}, \xi_{\mathrm{u}}, \xi_{\mathrm{w}}) \subseteq evol(\xi_{\mathrm{x}}, \xi_{\mathrm{u}}, \xi'_{\mathrm{w}})$. Thus, an increase of the uncertainty of the model may translate into a widening of the range of the possible behaviours of the CPS.

Proposition 8 (Monotonicity vs. uncertainty)**.** *Let $M = E \bowtie P$ with $E = \langle \xi_{\mathrm{x}}, \xi_{\mathrm{u}}, \xi_{\mathrm{w}}, evol, \xi_{\mathrm{e}}, meas, inv\rangle$. If $\xi_{\mathrm{w}} \le \xi'_{\mathrm{w}}$ and $M \xrightarrow{t} M'$, then $M\{^{\xi'_{\mathrm{w}}}/_{\xi_{\mathrm{w}}}\} \xrightarrow{t} M'\{^{\xi'_{\mathrm{w}}}/_{\xi_{\mathrm{w}}}\}$.*

However, a wider uncertainty in the model doesn't always correspond to a widening of the possible behaviours of the CPS. In fact, this depends on the *intrinsic tolerance* of a CPS with respect to changes in the uncertainty function.

4. Note that, apart from deadlocks, the tolerance of a CPS to an attack strongly depends on the capability of its cyber component to recognise abnormal physical states.

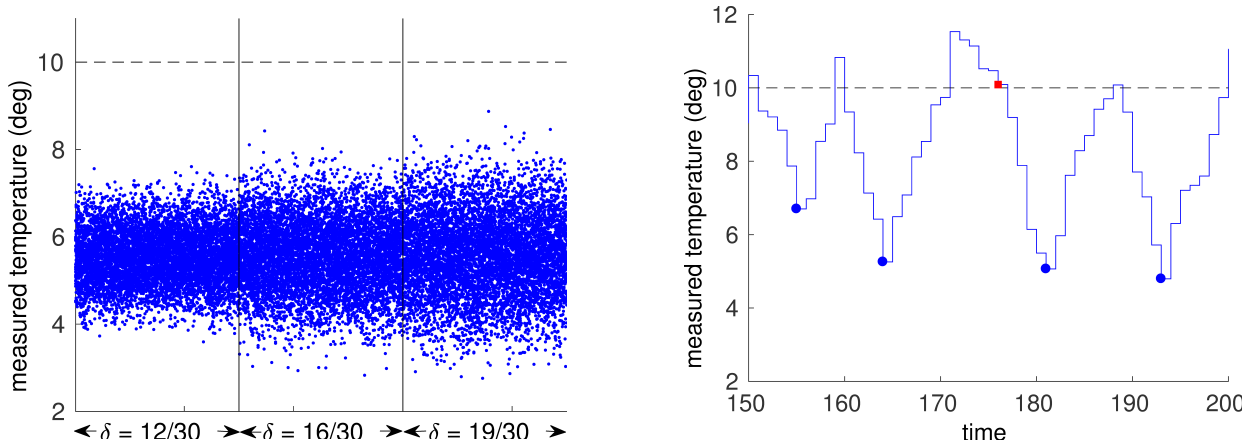

Figure 3. Simulations on the CPS of Example 8

Definition 11 (System ξ-tolerance). *Let* $E = \langle \xi_{\mathrm{x}}, \xi_{\mathrm{u}}, \xi_{\mathrm{w}}, evol, \xi_{\mathrm{e}}, meas, inv \rangle$ *with* $\xi_{\mathrm{w}} \in \mathbb{R}^{\hat{\mathcal{X}}}$. *A CPS* $M = E \bowtie P$ *is* ξ-tolerant, *for* $\xi \in \mathbb{R}^{\hat{\mathcal{X}}}$ *and* $\xi \geq 0$, *if*

$$\xi = \sup \{ \xi' : M \simeq M\{^{\xi_{\mathrm{w}}+\eta}/_{\xi_{\mathrm{w}}}\} \text{ for any } 0 \leq \eta \leq \xi' \}.$$

Intuitively, if a CPS M has been designed with a given uncertainty ξ_{w}, but M is actually ξ-tolerant, with $\xi > 0$, then the uncertainty ξ_{w} is somehow underestimated: the real uncertainty of M is given by $\xi_{\mathrm{w}}+\xi$. This information is quite important when trying to estimate the impact of an attack on a CPS. In fact, if a system M has been designed with a given uncertainty ξ_{w}, but M is actually ξ-tolerant, with $\xi > 0$, then a cyber-physical attack has (at least) a "room for manouvre" ξ to degrade the whole CPS without being observed (and hence detected). Thus, in general, the tolerance ξ should be as small as possible.

Example 8. *The CPS Sys of Example 2 is* $\frac{7}{30}$*-tolerant:*

$$\frac{7}{30} = \sup \{ \xi' : Sys \simeq Sys\{^{\delta+\eta}/_{\delta}\} \text{ for any } 0 \leq \eta \leq \xi' \} .$$

Since $\delta + \xi = \frac{12}{30} + \frac{7}{30} = \frac{19}{30}$, *the proof of this statement relies on the following proposition.*

Proposition 9. *Let Sys be the CPS defined in Example 2. Then:*

- $Sys \simeq Sys\{^{\gamma}/_{\delta}\}$, *for* $\gamma \in (\frac{12}{30}, \frac{19}{30})$,
- $Sys \not\simeq Sys\{^{\gamma}/_{\delta}\}$, *for* $\gamma > \frac{19}{30}$.

In Figure 3, the left graphic collects three pools of 300 simulations each, showing that the range of the temperatures measured after 5 time units of cooling becomes wider when increasing the uncertainty δ of the CPS *Sys* of Example 2. However, in any case, these temperatures are always below 10 and no alarm is fired. The right graphic shows an evolution of *Sys* with an uncertainty of $\frac{29}{30}$ degrees: the red box denotes a state in which the alarm is fired because the cooling cycle wasn't sufficient to drop the temperature below 10 (in that case, the controller imposed 5 further rounds of cooling).

Before giving a formalisation of our metric, we recall that when we defined the uncertainty function, we assumed operators for comparing the accuracy of two CPSs. Thus, for instance, for $\xi_{\mathrm{w}}, \xi'_{\mathrm{w}} \in \mathbb{R}^{\hat{\mathcal{X}}}$, we write $\xi_{\mathrm{w}} \leq \xi'_{\mathrm{w}}$ if $\xi_{\mathrm{w}}(x) \leq \xi'_{\mathrm{w}}(x)$, for any $x \in \hat{\mathcal{X}}$. Similarly, we write $\xi_{\mathrm{w}} + \xi'_{\mathrm{w}}$ to denote a function $\xi''_{\mathrm{w}} \in \mathbb{R}^{\hat{\mathcal{X}}}$ such that $\xi''_{\mathrm{w}}(x) = \xi_{\mathrm{w}}(x) + \xi'_{\mathrm{w}}(x)$, for any $x \in \hat{\mathcal{X}}$.

Now, everything is in place to provide our metric to estimate the impact of an attack.

Definition 12 (Impact of attack). *Let* $M = E \bowtie_{\mathcal{S}} P$ *be an honest CPS with* $E = \langle \xi_{\mathrm{x}}, \xi_{\mathrm{u}}, \xi_{\mathrm{w}}, evol, \xi_{\mathrm{e}}, meas, inv \rangle$ *and* $\xi_{\mathrm{w}} \in \mathbb{R}^{\hat{\mathcal{X}}}$. *We say that an attack* A *has* definitive impact ξ *to the system* M *if*

$$\xi = \inf \{ \xi' : \xi' \in \mathbb{R}^{\hat{\mathcal{X}}} \wedge \xi' > 0 \wedge M \parallel A \sqsubseteq M\{^{\xi_{\mathrm{w}}+\xi'}/_{\xi_{\mathrm{w}}}\} \}.$$

We say that an attack A *has* at least impact ξ *to the system* M *if there is a time* $m \in \mathbb{N}^+$ *such that*

$$\xi = \inf \{ \xi' : \xi' \in \mathbb{R}^{\hat{\mathcal{X}}} \wedge \xi' > 0 \wedge M \parallel A \sqsubseteq_{[m,n]} M\{^{\xi_{\mathrm{w}}+\xi'}/_{\xi_{\mathrm{w}}}\} \text{ for some } n \in \mathbb{N} \cup \infty \}.$$

Intuitively, with this definition, we can establish either the maximum impact of the attack A on system M or a lower bound of the impact at a specific time m. In the latter case, by definition of $\sqsubseteq_{[m,n]}$, there are two possibilities: either the impact of the attack keeps growing after time m, or in the time interval $m+1$, the system under attack deadlocks.

Example 9. *Let us consider the attack* A_n *of Example 6, for* $n = 10$. *Then,* A_{10} *has a definitive impact of* $\frac{7}{30}$ *to the CPS Sys defined in Example 2. (Sys has only one state variable, so we can assume* ξ_{w} *to be a constant). Formally,*

$$\frac{7}{30} = \inf \{ \xi' : \xi' > 0 \wedge Sys \parallel A_{10} \sqsubseteq Sys\{^{\delta+\xi'}/_{\delta}\} \} .$$

Here, the attack can influence the system by causing unexpected alarms. Since $\delta + \xi = \frac{12}{30} + \frac{7}{30} = \frac{19}{30}$, *the proof of this statement relies on the following proposition.*

Proposition 10. *Let Sys be the CPS defined in Example 2, and* A_{10} *be the attack defined in Example 9. Then:*

- $Sys \parallel A_{10} \not\sqsubseteq Sys\{^{\gamma}/_{\delta}\}$, *for* $\gamma \in (\frac{12}{30}, \frac{19}{30})$,
- $Sys \parallel A_{10} \sqsubseteq Sys\{^{\gamma}/_{\delta}\}$, *for* $\gamma > \frac{19}{30}$.

The following example proposes a different attack with a different impact on Sys.

Example 10. *Let us consider the attack A of Example 4. Then, A has a definitive impact of 0.58 to the CPS Sys defined in Example 2 . Formally:*

$$0.58 = \inf\left\{\, \xi' : \xi' > 0 \ \wedge\ Sys \parallel A \sqsubseteq Sys\{^{\delta+\xi'}/_{\delta}\}\right\} .$$

Here, the attack can prevent the activation of the cooling system. As a consequence, the temperature will keep growing until the CPS deadlocks. Since $\delta+\xi = 0.4+0.58 = 0.98$, the proof of this statement relies on the following proposition.

Proposition 11. *Let Sys be the CPS defined in Example 2, and A be the attack defined in Example 10. Then:*

- *$Sys \parallel A \not\sqsubseteq Sys\{^{\gamma}/_{\delta}\}$, for $\gamma \in (0.4, 0.98)$,*
- *$Sys \parallel A \sqsubseteq Sys\{^{\gamma}/_{\delta}\}$, for $\gamma > 0.98$.*

In Definition 12, we have provided an instrument to estimate the impact of a successful attack. However, there is at least another question that might arise in the mind of a CPS designer: "Is there a way to estimate the chances that an attack will be successful during the execution of my CPS"? To paraphrase in a more operational manner: how many execution traces of my system are prone to be attacked by a specific attack?

For instance, consider again the simple attack A_m of class $\{\lightning cool?, \lightning cool!\}[m, 1]$, proposed in Example 3:

$$\mathsf{idle}^{m-1}.\mathsf{timeout}\lfloor \mathsf{read}\,\lightning cool(x).\mathsf{if}\ (x = \mathsf{off})\ \{\mathsf{write}\,\lightning cool\langle \mathsf{off}\rangle\}\rfloor .$$

Here, in the m-th time slot the attack tries to eavesdrop a command used by the controller to turn on the cooling. The attack is very quick and condensed in a single time slot. The question is: what are the chances of success of such a quick attack?

Figure 4 provides a graphical representation of an experiment where we launched 10000 executions of our CPS in isolation, lasting 700 time units each. From the aggregated data contained in this graphic, we note that after a transitory phase (whose length depends on several things: the uncertainty δ, the initial state of the system, the length of the cooling activity, etc.) that lasts around 300 time slots, the rate of success of the attack A_m is around 10%. That is, after the transitory phase, choosing an appropriate $m > 300$, around 10% of the execution traces of the CPS Sys are prone to the attack A_m. This is because, in average, 10% of the commands to turn on the cooling system are stolen by the attack causing an alarm. The reader may wonder why exactly the 10%. This depends on the periodicity of our CPS, where in average the cooling is activated every 10 time slots. For instance, reducing the time of cooling to 4 time units, the number of activations would be, in average, every 8 time slots, thus the rate of success of this attack would be 12.5%. This example shows that, as pointed out in [9], the effectiveness of a cyber-physical attack depends considerably on the information the attack has about the functionality of the whole CPS.

For instance, if the attacker were not aware of the exact periodicity of the CPS, she might try, if possible, to repeat

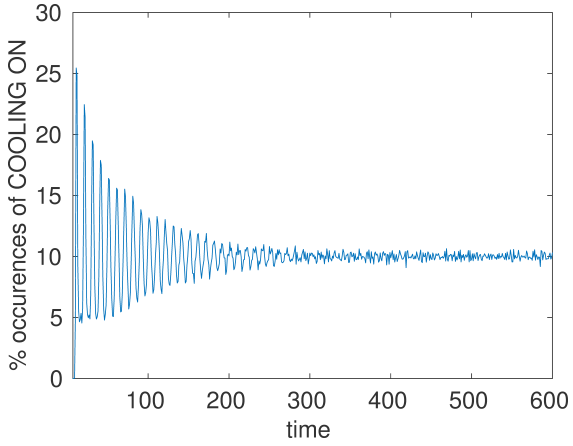

Figure 4. A quantitative analysis of the attack of Example 3.

the attack on more consecutive time slots. In this case, the left graphic of Figure 5 says that the rate of success of the attack increases linearly with the length of the attack itself.[5] Thus, if the attack of Example 3 were iterated for 10 consecutive time slots, say

$$\begin{array}{rcl} A_m^{10} & = & \mathsf{idle}^{m-1}.B_{10} \\ B_i & = & \mathsf{timeout}\lfloor \mathsf{read}\,\lightning cool(x).\mathsf{if}\ (x = \mathsf{off}) \\ & & \{\mathsf{write}\,\lightning cool\langle \mathsf{off}\rangle.\mathsf{idle}.B_{i-1}\}\rfloor B_{i-1},\ \text{for } 1{\leq}i{\leq}10 \\ B_0 & = & \mathsf{nil} \end{array}$$

becoming an attack of class $\{\lightning cool?, \lightning cool!\}[m, 10]$, the rate of success would be almost 100%. This is, of course, because our running example is very simple. In more sophisticated CPSs, the same analysis would be much harder [8].

Finally, consider the following generalisation of the attack of Example 6:

$$\begin{array}{rcl} A_0^k & = & \mathsf{nil} \\ A_n^k & = & \mathsf{timeout}\lfloor \mathsf{read}\,\lightning s_{\mathrm{t}}(x). \\ & & \mathsf{timeout}\lfloor \mathsf{write}\,\lightning s_{\mathrm{t}}\langle x{-}k\rangle.\mathsf{idle}.A_{n-1}^k\rfloor A_{n-1}^k\rfloor A_{n-1}^k\,, \end{array}$$

for $1 \leq n \leq 15$ and $2 \leq k \leq 10$. Here, the attack decreases the sensed temperature of an offset k. Now, suppose to launch this attack after, say, 300 time slots (i.e., after the transitory phase). Formally, we define the attack:

$$B_n^k \quad = \quad \mathsf{idle}^{300}.A_n^k$$

of class $\{\lightning s_{\mathrm{t}}!, \lightning s_{\mathrm{t}}?\}[301, n]$. In this case, the right graphic of Figure 5 provides a graphical representation of the percentage of alarms on 5000 execution traces lasting 100 time units each. Thus, for instance, an attack lasting $n = 8$ time units with an offset $k = 5$ affects around 40% of the execution traces of the CPS.

5. Conclusions and related work

The goal of this paper was to apply *formal methods* to lay and streamline *theoretical foundations* to reason about and

5. Data obtained by attacking the CPS after the transitory period.

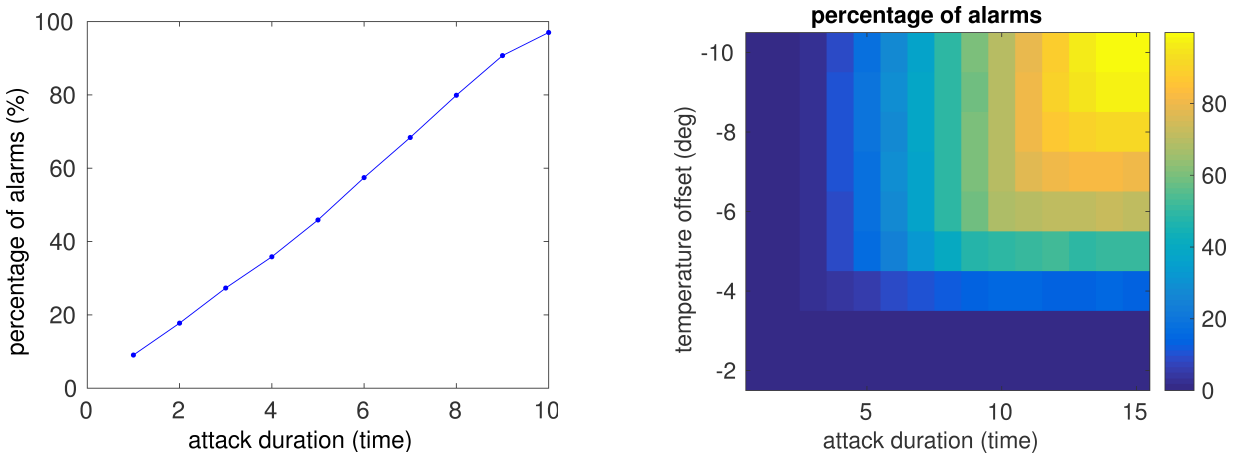

Figure 5. A quantitative analysis of two different attacks.

detect attacks to physical devices of CPSs. To that end, we have proposed a hybrid process calculus, called CCPSA, as a formal *specification language* to model physical and cyber components of CPSs as well as cyber-physical attacks. Based on CCPSA and its labelled transition semantics, we have formalised a threat model for CPSs by grouping attacks in classes (Definition 9), according to the target physical devices and two timing parameters: begin and duration of the attacks. Then, we relied on the behavioural semantics of CCPSA, trace equivalence, to *assess attack tolerance/vulnerability with respect to a given attack* (Definition 10). Such a tolerance may hold ad-infinitum or for a limited amount of time. In the latter case, the CPSs is vulnerable and the attack affects the observable behaviour of the system only after a certain point in time, when the attack itself may already be achieved (in this case, we call it stealthy attack) or still working. Here, a crucial decision for the CPS-designer is the choice of which devices are worth to protect with cryptographic means to increase the tolerance of the CPS to certain attacks (Example 5). However, making a CPS more robust with respect to a certain attack is a delicate task, as changes in the cyber component may open the way to new attacks (Example 7). Finally, we provided a metric to estimate the impact of a successful attack on a CPS (Definition 12) together with possible quantifications of the success chances of an attack. We have illustrated our concepts by means of a running example, focusing in particular on a formal treatment of both integrity and DoS attacks to sensors and actuators of CPSs. As we remarked above, our example is simple but far from trivial and designed to describe a wide number of attacks (and a more complex example wouldn't have been more instructive).

For a comprehensive up-to-date survey on CPS security, we refer the reader to [28]. Here, we would like to point out that among the 118 papers analysed in [28], 50 adopt a discrete notion of time similar to ours, 13 a continuous one, 48 a quasi-static time model, and the remaining papers use a hybrid time model. Most of these papers investigate attacks on sensors and actuators and their protection, by relying on *simulation test systems* to validate the results. While simulation tools (such as *Matpower* [30], used in the power system application domain) are widely used in the context of CPSs, the complexity of CPSs would call for (semi-)automated tools to provide both an exhaustive search of all possible behaviours of the system, and exact, rather than approximate, quantitative results.

In the remainder of this section, we will focus on the papers that have been influential for us and/or are most closely related to our work.

Huang et al. [12] were among the first to propose threat models for CPSs: they proposed models for integrity and DoS attacks, and evaluated the physical and economic consequences of the attacks on a chemical reactor system. They, along with Krotofil et al. [13], [14], have also stressed the important role played by timing parameters on both integrity and DoS attacks. For instance, the Tennessee Eastman process control challenge problem is used in [14] to investigate the problem of an attacker who has to identify (in real time) the optimal moment to launch a DoS attack.

Gollmann et al. [9] discussed possible goals of cyber-physical attack: *equipment damagage*, *production damage*, and *compliance violation*. Furthermore, they fixed the *stages* that a cyber-physical attack should go through before achieving her goals: *access*, *discovery*, *control*, *damage*, and *cleanup*. The stages of this approach are demonstrated via a case study, a simulation of a vinyl acetate monomer plant. In our paper, we only focus on the fourth stage, damage, where the attack has already a rough idea of the control plan of the target CPS.

Genge et al. [8] showed how different attacks may have different impact on CPSs, and proposed a methodology for assessing the impact of attacks on critical infrastructures, based on system dynamics and sensitivity analysis.

Teixeira et al. [24] proposed a three-dimensional attack space, for both cyber and physical attacks. The dimensions of the threat model are: the attacker's a priori system model knowledge, disclosure, and disruption resources. This allowed them to analyse DoS, replay, zero-dynamics, and

bias injection attacks. Furthermore, the attack policy for each scenario is described and the attack's impact is characterised using the concept of safe sets.

A number of approaches have been proposed for modelling CPSs using formal methods. For instance, *hybrid automata* [2] combine finite state transition systems with discrete variables (whose values capture the state of the modelled discrete or cyber components) and continuous variables (whose values capture the state of the modelled continuous or physical components). *Hybrid process algebras* [5] are a powerful tool for reasoning about physical systems and provide techniques for analysing and verifying protocols for hybrid automata. CCPSA shares some similarities with the ϕ-calculus [21], an hybrid extension of the π-calculus [17]. In the ϕ-calculus, an hybrid system is a pair (E, P), where E is the environment and P is the process interacting with the environment. Unlike CCPSA, in the ϕ-calculus, given a system (E, P), the process P can dynamically change the evolution law in E. The ϕ-calculus does not have a representation of physical devices and measurement law, which are instead crucial for us to model cyber-physical attacks that operate in a timely fashion on sensors and actuators.

We are aware of a (still fairly small) number of works that use formal methods for CPS security, although they apply methods, and most of the time have goals, that are quite different from ours. We have space to discuss only a few of them.

Burmester et al. [3] proposed a framework for modelling the security of CPSs based on *hybrid timed automa* [11] in which the behaviour of the adversary is controlled by a threat model that captures the cyber aspects (with discrete values) and the physical aspects (with continuous values) of CPSs. The threat framework is based on the traditional Byzantine faults model for cryptographic security in which the basic security features and requirements, as specified by the security policies, are used to identify system vulnerabilities. However, as remarked in [24], cyber-physical attacks and faults have inherently distinct characteristics, which pose different challenges on the secure control and fault-tolerant approaches. Faults are considered as physical events that affect the system behaviour, where simultaneous events are assumed to be non-colluding, i.e., the events do not act in a coordinated way. On the other hand, as seen in Example 5, cyber attacks may be performed over a significant number of attack points and in a coordinated fashion.

Cómbita et al. [4] and Zhu and Basar [29] applied *game theory* to capture the conflict of goals between an attack who seeks to maximise the damage inflicted to a CPS's security and a defender who aims to minimise it [16].

In [26], Vigo presented an attack scenario that addresses some of the peculiarities of a cyber-physical adversary, and discussed how this scenario relates to other attack models popular in the security protocol literature. In [27] (and related papers), he then proposed an untimed calculus of broadcasting processes instrumented with a theory that allows for modelling and reasoning about cryptographic primitives, and equipped with explicit notions of communication failure and unwanted communication. These works differ quite considerably from ours, e.g., they focus on DoS attacks without taking into consideration timing aspects or their impact on the system under attack.

Finally, we are aware of three very recent papers that were developed in parallel to ours. In [19], Rocchetto and Tippenhaur introduce a taxonomy to classify the diverse attacker models that have been proposed in the CPS security community, also outlining requirements for generalised attacker models for CPSs. In [20], the same authors propose an extended Dolev-Yao attacker model suitable for arguments about CPS security. In [18], Nigam et al. extend previous works on Timed Intruder Models for Cyber-Physical Security Protocols by bounding the number of such intruders required for the automated verification of such protocols. We leave a detailed comparison of our approach with these ones as future work.

While much work is of course still to be done, we believe that the definitions, results and examples given in this paper will provide a stepping stone for the development of formal and automated tools to analyse security properties of CPSs. To that end, we will consider applying, possibly after proper enhancements, existing tools and frameworks for automated security protocol analysis, resorting to the development of an dedicated tool if existing ones prove not up to the task.

In parallel to this, we will consider further security properties and concrete examples of CPSs, as well as other kinds of cyber-physical attackers and attacks, e.g., periodic attacks. This will allow us to refine the classes of attacks we have given here (e.g., by formalising a type system that could provide the basis for a static analysis), as well as provide a formal definition of when a CPS is more secure than another and ultimately be able to design, by progressive refinement, secure variants of a vulnerable CPSs.

We also aim to extend the preliminary quantitative analysis we have given here by developing a suitable behavioural theory ensuring that our trace semantics considers also the probability of a trace to actually occur. We expect that the discrete time stochastic hybrid systems of [1] will be useful to that extent.

References

[1] A. Abate, S. Amin, M. Prandini, J. Lygeros, and S. Sastry. Probabilistic reachability and safe sets computation for discrete time stochastic hybrid systems. In *CDC*, pages 258–263. IEEE, 2006.

[2] R. Alur, C. Courcoubetis, T. Henzinger, and P. Ho. Hybrid automata: An algorithmic approach to the specification and verification of hybrid systems. In *Hybrid Systems*, pages 209–229. Springer, 1993.

[3] M. Burmester, E. Magkos, and V. Chrissikopoulos. Modeling security in cyber-physical systems. *IJCIP*, 5(3-4):118–126, 2012.

[4] L. F. Cómbita, J. Giraldo, A. A. Cárdenas, and N. Quijano. Response and reconfiguration of cyber-physical control systems: A survey. In *Automatic Control 2015*, pages 1–6. IEEE, 2015.

[5] P. Cuijpers and M. Reniers. Hybrid process algebra. *JLAP*, 62(2):191–245, 2005.

[6] N. Falliere, L. Murchu, and E. Chien. W32.Stuxnet Dossier, 2011.

[7] R. Focardi and F. Martinelli. A uniform approach for the definition of security properties. In *Symposium on Formal Methods*, volume I, pages 794–813, 1999.

[8] B. Genge, I. Kiss, and P. Haller. A system dynamics approach for assessing the impact of cyber attacks on critical infrastructures. *Int. J. Critical Infrastructure Protection*, 10:3–17, 2015.

[9] D. Gollmann, P. Gurikov, A. Isakov, M. Krotofil, J. Larsen, and A. Winnicki. Cyber-Physical Systems Security: Experimental Analysis of a Vinyl Acetate Monomer Plant. In *CPSS*, pages 1–12. ACM, 2015.

[10] M. Hennessy and T. Regan. A process algebra for timed systems. *I&C*, 117(2):221–239, 1995.

[11] T. A. Henzinger. The theory of hybrid automata. In *LICS*, pages 278–292. IEEE Computer Society, 1996.

[12] Y. Huang, A. A. Cárdenas, S. Amin, Z. Lin, H. Tsai, and S. Sastry. Understanding the physical and economic consequences of attacks on control systems. *IJCIP*, 2(3):73–83, 2009.

[13] M. Krotofil and A. A. Cárdenas. Resilience of Process Control Systems to Cyber-Physical Attacks. In *NordSec*, LNCS 8208, pages 166–182. Springer, 2013.

[14] M. Krotofil, A. A. Cárdenas, J. Larsen, and D. Gollmann. Vulnerabilities of cyber-physical systems to stale data - Determining the optimal time to launch attacks. *Int. J. Critical Infrastructure Protection*, 7(4):213–232, 2014.

[15] E. Levy. Crossover: Online pests plaguing the offline world. *IEEE Security & Privacy*, 1(6):71–73, 2003.

[16] M. Manshaei, Q. Zhu, T. Alpcan, T. Basar, and J.-P. Hubaux. Game theory meets network security and privacy. *ACM Comput. Surv.*, 45(3):25, 2013.

[17] R. Milner. The polyadic π-calculus: a tutorial. Technical Report ECS–LFCS–91–180, LFCS, 1991.

[18] V. Nigam, C. Talcott, and A. A. Urquiza. Towards the Automated Verification of Cyber-Physical Security Protocols: Bounding the Number of Timed Intruders. In *Esorics*, LNCS 9878-9. Springer, 2016.

[19] M. Rocchetto and N. O. Tippenhauer. On Attacker Models and Profiles for Cyber-Physical Systems. In *Esorics*, LNCS 9878-9. Springer, 2016.

[20] M. Rocchetto and N. O. Tippenhauer. CPDY: Extending the Dolev-Yao Attacker with Physical-Layer Interactions. In *ICFEM 2016*, to appear.

[21] W. C. Rounds and H. Song. The ϕ-calculus: A language for distributed control of reconfigurable embedded systems. In *HSCC*, pages 435–449. Springer, 2003.

[22] S. Sicari, A. Rizzardi, L. Grieco, and A. Coen-Porisini. Security, privacy and trust in internet of things: The road ahead. *Computer Networks*, 76:146–164, 2015.

[23] J. Slay and M. Miller. Lessons Learned from the Maroochy Water Breach. In *Critical Infrastructure Protection*, IFIP 253, pages 73–82. Springer, 2007.

[24] A. Teixeira, I. Shames, J. Sandberg, and K. H. Johansson. A secure control framework for resource-limited adversaries. *Automatica*, 51:135–148, 2015.

[25] U.S. Chemical Safety and Hazard Investigation Board, T2 Laboratories Inc. Reactive Chemical Explosion: Final Investigation Report. Report No. 2008-3-I-FL, 2009.

[26] R. Vigo. The Cyber-Physical Attacker. In *SAFECOMP*, LNCS 7613, pages 347–356. Springer, 2012.

[27] R. Vigo. *Availability by Design: A Complementary Approach to Denial-of-Service*. PhD thesis, DTU Compute, Danish Technical University, 2015.

[28] Y. Zacchia Lun, A. D'Innocenzo, I. Malavolta, and M. D. Di Benedetto. Cyber-Physical Systems Security: a Systematic Mapping Study. *CoRR*, abs/1605.09641, 2016.

[29] Q. Zhu and T. Basar. Game-theoretic methods for robustness, security, and resilience of cyberphysical control systems: games-in-games principle for optimal cross-layer resilient control systems. *IEEE Control Systems*, 35(1):46–65, 2015.

[30] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas. MATPOWER: Steady-State Operations, Planning and Analysis Tools for Power Systems Research and Education. *Power Systems, IEEE Transactions on*, 26:12–19, 2011.

Appendix

1. Proof of § 2

In order to prove Proposition 1 and Proposition 2 we use the following lemma that formalises the invariant properties binding the state variable $temp$ with the activity of the cooling system.

Intuitively, when the cooling system is inactive the value of the state variable $temp$ lays in the interval $[0, 11.5]$. Furthermore, if the coolant is not active and the variable $temp$ lays in the interval $(10.1, 11.5]$ then the cooling will be turned on in the next time slot. Finally, when active then cooling system will remain so for $k \in 1..5$ time slots (counting also the current time slot) being the variable $temp$ in the real interval $(9.9 - k*(1+\delta), 11.5 - k*(1-\delta)]$.

Lemma 1. *Let Sys be the system defined in Example 2. Let*

$$Sys = Sys_1 \xrightarrow{t_1} \xrightarrow{\mathsf{idle}} Sys_2 \xrightarrow{t_2} \xrightarrow{\mathsf{idle}} \cdots \xrightarrow{t_{n-1}} \xrightarrow{\mathsf{idle}} Sys_n$$

such that the traces t_j contain no idle*-actions, for any $j \in 1..n-1$, and for any $i \in 1..n$ $Sys_i = E_i \bowtie P_i$ with $E_i = \langle \xi^i_{\mathrm{x}}, \xi^i_{\mathrm{u}}, \delta, evol, \epsilon, meas, inv \rangle$. Then, for any $i \in 1..n-1$ we have the following:*

1) *if $\xi^i_{\mathrm{u}}(cool) = \mathsf{off}$ then $\xi^i_{\mathrm{x}}(temp) \in [0, 11.1 + \delta]$;*
2) *if $\xi^i_{\mathrm{u}}(cool) = \mathsf{off}$ and $\xi^i_{\mathrm{x}}(temp) \in (10.1, 11.1 + \delta]$ then, in the next time slot, $\xi^{i+1}_{\mathrm{u}}(cool) = \mathsf{on}$;*
3) *if $\xi^i_{\mathrm{u}}(cool) = \mathsf{on}$ then $\xi^i_{\mathrm{x}}(temp) \in (9.9 - k*(1+\delta), 11.1 + \delta - k*(1-\delta)]$, for some $k \in 1..5$ such that $\xi^{i-k}_{\mathrm{u}}(cool) = \mathsf{off}$ and $\xi^{i-j}_{\mathrm{u}}(cool) = \mathsf{on}$, for $j \in 0..k-1$.*

Proof. Let us denote with v_i the values of the state variable $temp$ in the systems Sys_i, i.e., $\xi^i_{\mathrm{x}}(temp) = v_i$. Moreover we will say that the coolant is active (resp., is not active) in Sys_i if $\xi^i_{\mathrm{u}}(cool) = \mathsf{on}$ (respectively, $\xi^i_{\mathrm{u}}(cool) = \mathsf{off}$).

The proof is by mathematical induction on n, i.e., the number idle-actions of our traces.

The case base $n = 1$ follows directly from the definition of Sys.

Let prove the inductive case. We assume that the three statements holds for $n - 1$ and we prove that they also hold for n.

1) Let us assume that the cooling is not active in Sys_n, then we prove that $v_n \in [0, 11.1 + \delta]$.
   We consider separately the cases in which the coolant is active or not in Sys_{n-1}
   - Suppose the coolant is not active in Sys_{n-1} (and inactive in Sys_n).
     By inductive hypothesis we have $v_{n-1} \in [0, 11.1+\delta]$. Furthermore, if $v_{n-1} \in (10.1, 11.1 + \delta]$ then, by inductive hypotheses, the coolant must be active in Sys_n. Since we know in Sys_n the cooling is not active it follows that $v_{n-1} \in [0, 10.1]$. Furthermore, Sys_n the temperature will increase of a value laying in the interval $[1-\delta, 1+\delta] = [0.6, 1.4]$. Thus v_n will be in $[0.6, 11.1 + \delta] \subseteq [0, 11.1 + \delta]$.
   - Suppose the coolant is active in Sys_{n-1} (and inactive in Sys_n).
     By inductive hypothesis $v_{n-1} \in (9.9 - k * (1 + \delta), 11.1 + \delta - k*(1-\delta)]$ for some $k \in 1..5$ such that the coolant is not active in Sys_{n-1-k} and is active in $Sys_{n-k}, \ldots, Sys_{n-1}$.
     The case $k \in \{1, \ldots, 4\}$ is not admissible. In fact if $k \in \{1, \ldots, 4\}$ then the coolant would be active for less than 5 idle-actions as we know that Sys_n is inactive. Hence, it must be $k = 5$. Since $\delta = 0.4$ and $k = 5$, it holds that $v_{n-1} \in (9.9-5*1.4, 11.1+0.4-5*0.6] = (2.8, 8.6]$. Moreover, since the coolant is active for 5 time intervals, in Sys_{n-1} the controller and the IDS synchronise together via channel $sync$ and hence the IDS checks the temperature. Since $v_{n-1} \in (2.8, 8.6]$ the IDS process sends to the controller a command to stop the cooling, and the controller will switch off the cooling system. Thus, in the next time slot, the temperature will increase of a value laying in the interval $[1 - \delta, 1 + \delta] = [0.6, 1.4]$. As a consequence, in Sys_n we will have $v_n \in [2.8+0, 6, 8.6+1.4] = [3.4, 10] \subseteq [0, 11.1+\delta]$.
2) Let us assume that the coolant is not active in Sys_n and $v_n \in (10.1, 11.1 + \delta]$, then we prove that the coolant is active in Sys_{n+1}. Since the coolant is not active in Sys_n then it will check the temperature before the next time slot. Since $v_n \in (10.1, 11.1 + \delta]$ and $\epsilon = 0.1$, then the process $Ctrl$ will sense a temperature greater than 10 and the coolant will be turned on. Thus the coolant will be active in Sys_{n+1}.
3) Let us assume that the coolant is active in Sys_n, then we prove that $v_n \in (9.9-k*(1+\delta), 11.1+\delta-k*(1-\delta)]$ for some $k \in 1..5$ and the coolant is not active in Sys_{n-k} and active in $Sys_{n-k+1}, \ldots, Sys_n$.
   We separate the case in which the coolant is active in Sys_{n-1} from that in which is not active.
   - Suppose the coolant is not active in Sys_{n-1} (and active in Sys_n).
     In this case $k = 1$ as the coolant is not active in Sys_{n-1} and it is active in Sys_n. Since $k = 1$, we have to prove $v_n \in (9.9-(1+\delta), 11.1+\delta-(1-\delta)]$. However, since the coolant is not active in Sys_{n-1} and is active in Sys_n it means that the coolant has been switched on in Sys_{n-1} because the sensed temperature was above 10 (since $\epsilon = 0.1$ this may happen only if $v_{n-1} > 9.9$). By inductive hypothesis, since the coolant is not active in Sys_{n-1}, we have that $v_{n-1} \in [0, 11.1 + \delta]$. Therefore, from $v_{n-1} > 9.9$ and $v_{n-1} \in [0, 11.1 + \delta]$ it follows that $v_{n-1} \in (9.9, 11.1 + \delta]$. Furthermore, since the coolant is active in Sys_n, the temperature will decrease of a value in $[1 - \delta, 1 + \delta]$ and therefore $v_n \in (9.9-(1+\delta), 11.1+\delta-(1-\delta)]$ which concludes this case of the proof.
   - Suppose the coolant is active in Sys_{n-1} (and active in Sys_n as well).
     By inductive hypothesis there is $h \in 1..5$ such that

$v_{n-1} \in (9.9 - h*(1+\delta), 11.1 + \delta - h*(1-\delta)]$ and the coolant is not active in Sys_{n-1-h} and is active in $Sys_{n-h}, \ldots, Sys_{n-1}$.
The case $h = 5$ is not admissible. In fact, since $\delta = 0.4$, if $h = 5$ then $v_{n-1} \in (9.9 - 5*1.4, 11.1 + \delta - 5*0.6] = (2.8, 8.6]$. Furthermore, since the cooling system is active since 5 time instants, in Sys_{n-1} the controller and the IDS synchronise together via channel $sync$, and the IDS checks the received temperature. As $v_{n-1} \in (2.8, 8.6]$ then the IDS sends to the controller via channel ins the command stop. This implies that the controller should turn off the cooling system. In contradiction with the the hypothesis that the coolant is active in Sys_n.
Hence it must be $h \in 1..4$. Let us prove that for $k = h + 1$ we obtain our result. Namely we have to prove that, for $k = h + 1$, (i) $v_n \in (9.9 - k*(1+\delta), 11.1 + \delta - k*(1-\delta))$, and (ii) the coolant is not active in Sys_{n-k} and active in $Sys_{n-k+1}, \ldots, Sys_n$.
Let us prove the statement (i). By inductive hypotheses, it holds that $v_{n-1} \in (9.9 - h*(1+\delta), 11.1 + \delta - h*(1-\delta))$. Since the coolant is active in Sys_n then the temperature will decrease Hence, $v_n \in (9.9 - (h+1)*(1+\delta), 11.1 + \delta - (h+1)*(1-\delta))$. Therefore, since $k = h + 1$, we have that $v_n \in (9.9 - k*(1+\delta), 11.1 + \delta - k*(1-\delta))$.
Let us prove the statement (ii). By inductive hypothesis the coolant is inactive in Sys_{n-1-h} and it is active in $Sys_{n-h}, \ldots, Sys_{n-1}$. Now, since the coolant is active in Sys_n, for $k = h + 1$, we have that the coolant is not active in Sys_{n-k} and is active in $Sys_{n-k+1}, \ldots, Sys_n$ which concludes this case of the proof.

□

Proof of Proposition 1. By Lemma 1 and since $\delta = 0.4$, the value of the state variable $temp$ is always in the real interval $[0, 11.5]$. As a consequence, the invariant of the system is never violated and the system never deadlocks. Moreover, after 5 idle -actions of cooling the state variable $temp$ is always in the real interval $(9.9 - 5*1.4, 11.1 + 0.4 - 5*0.6] = (2.9, 8.5]$. Hence the process IDS will never transmit on the channel $alarm$. □

Proof of Proposition 2. Let us prove the two statements separately.

- Since $\epsilon = 0.1$, if process $Ctrl$ senses a temperature above 10 (and hence Sys turns on the cooling) then the value of the state variable $temp$ is greater than 9.9. By Lemma 1 the value of the state variable $temp$ is always less or equal than $11.1 + \delta$. Therefore, if $Ctrl$ senses a temperature above 10, then the value of the state variable $temp$ is in $(9.9, 11.1 + \delta]$.
- By Lemma 1 (third item) the coolant can be active for no more than 5 time slots. Hence, by Lemma 1, when Sys turns off the cooling system the state variable $temp$ ranges over $(9.9 - 5*(1+\delta), 11.1 + \delta - 5*(1-\delta))$.

□

2. Proofs of § 3

Proof of Proposition 3. By rule induction of why $E \bowtie P \xrightarrow{\tau} E' \bowtie P'$. The proofs relies on the fact that recursive processes are always time-guarded in CCPSA. This ensures that processes Q and P' are always different. □

Proof of Proposition 5. By induction on the length of the trace. □

Proof of Proposition 4. We distinguish the two cases, depending on m.

- Let $m \leq 8$. We recall the cooling system is activated only when the sensed temperature is above 10. Since $\epsilon = 0.1$, when this happens the state variable $temp$ must be at least 9.9. Note that after $m - 1 \leq 7$ idle-actions, when the attack tries to interact with the controller of the actuator $cool$, the variable $temp$ may reach at most $7*(1+\delta) = 7*1.4 = 9.8$ degrees. Thus, the cooling system will not be activated and the attack will not have any effect.
- Let $m > 8$. By Proposition 1, the system Sys in isolation may never deadlock and it may never emit an output on channel $alarm$. Thus, any execution trace of the system Sys consists of a sequence of τ-actions and idle-actions.
  As a consequence, in order to prove the equality above it is enough to show the following three facts:
  - the system $Sys \parallel A_m$ may not deadlock in the first $m + 4$ time slots;
  - the system $Sys \parallel A_n$ may not emit any output in the first $m + 4$ time slots;
  - there is a trace in which the system $Sys \parallel A_m$ emits a sequence of outputs along channel $alarm$, starting from the $(m+5)$-th time interval and until the invariant gets violated.

  The first two facts are easy to show as the attack may steal the command addressed to actuator $cool$ only in the m-th time slot. Thus, until the time slot m, the whole system behaves correctly. In particular, by Proposition 1 and Proposition 2, no alarms and deadlocks occurs, and the temperature lies in the expected ranges. Any observable effect of the attack A_m, i.e., deadlock or output on channel $alarm$, requires at least further 5 time slots to occur. Thus, none of these observable effects will occur in the first $m + 4$ time slots.
  Let us show the third fact, i.e., that there is a trace in which the system $Sys \parallel A_m$ emits an output on channel $alarm$ in the $(m + 5)$-th time slot.
  When the process $Ctrl$ senses a temperature above 10, by Proposition 2 it follows that the actual value of the variable $temp$ is in the real interval $(9.9, 11.5]$ (recall that $\delta = 0.4$). Thus, the first activation of coolant may occur in *any* time slot k, with $9 \leq k \leq 21$, depending on the velocity of the system evolution. In fact, the fastest activation occurs after $\lceil \frac{9.9}{1+\delta} \rceil = \lceil \frac{9.9}{1.4} \rceil = 8$ time steps, while the slowest one after $\lceil \frac{11.5}{1-\delta} \rceil = \lceil \frac{11.5}{0.6} \rceil = 20$ time steps (they correspond to the trajectories blue and

red, respectively, of the right graphic of Figure 2). In other words, for any k, with $9 \leq k \leq 21$, there is a trace of the system where the coolant is activated in the time slot k. Now, for any of those time slots k, there is a trace in which the following activations of the coolant starts at all time slots $k + 10i$, for any $i \in \mathbb{N}$. In fact, we can always consider the trace in which, during the cooling, the temperature is decreased of $1.4 = 1 + \delta$ degrees per time unit, for 5 consecutive time units, and in the following 5 time units the temperature is increased of the same value, reaching the initial state in which the sensed temperature is above 10 and the coolant must be activated again. As a consequence, for any time slot $k \in \mathbb{N}$, $k > 8$ there is a trace of the system Sys, and hence also of $Sys \parallel A_k$, in which at the time slot k the system sends the command $cool!\mathsf{on}$ to the actuator $cool$. Now, if the attack A_m is active in that time slot (i.e., $m = k > 8$) the command will be stolen by the attack and it will never reach the actuator. Without that dose of coolant the temperature of the system, which in that moment is at least 9.9, will continue to grow for the following 5 time slots of a value $v \in [1-\delta\,,\, 1+\delta] = [0.6, 1.4]$. After 5 time slots, in the time slot $m + 5$, the controller and the IDS synchronise via channel $sync$, the IDS will detect a temperature above 10, and it will fire the output on channel $alarm$ saying to process $Ctrl$ to keep cooling. But $Ctrl$ will not send again the command $cool!\mathsf{on}$. Hence, the temperature would continue to increase and the process IDS will send a sequence of $alarm$ until the invariant of the environment gets violated.

□

In order to prove Proposition 6 we need a couple of lemmas.

Lemma 2 is a variant of Lemma 1 where the CPS Sys runs in parallel with the attack A_n defined in Example 6. Here, due to the presence of the attack, the temperature is 2 degrees higher with respect to the system Sys in isolation.

Lemma 2. *Let Sys be the system defined in Example 2 and A_n be the attack of Example 6 Let*

$$Sys \parallel A_n = Sys_1 \xrightarrow{t_1} \xrightarrow{\mathsf{idle}} \dots Sys_{n-1} \xrightarrow{t_n} \xrightarrow{\mathsf{idle}} Sys_n$$

such that the traces t_j contain no idle*-actions, for any $j \in 1..n-1$, and for any $i \in 1..n$ $Sys_i = E_i \bowtie P_i$ with $E_i = \langle \xi^i_{\mathrm{x}}, \xi^i_{\mathrm{u}}, \delta, evol, \epsilon, meas, inv \rangle$. Then, for any $i \in 1..n-1$ we have the following:*

- *if $\xi^i_{\mathrm{u}}(cool) = \mathsf{off}$, then $\xi^i_{\mathrm{x}}(temp) \in [0, 11.1 + 2 + \delta]$;*
- *if $\xi^i_{\mathrm{u}}(cool) = \mathsf{off}$ and $\xi^i_{\mathrm{x}}(temp) \in (10.1+2, 11.1+2+\delta]$, then we have $\xi^{i+1}_{\mathrm{u}}(cool) = \mathsf{on}$;*
- *if $\xi^i_{\mathrm{u}}(cool) = \mathsf{on}$, then $\xi^i_{\mathrm{x}}(temp) \in (9.9 + 2 - k * (1 + \delta), 11.1+2+\delta-k*(1-\delta)]$, for some $k \in 1..5$, such that $\xi^{i-k}_{\mathrm{u}}(cool) = \mathsf{off}$ and $\xi^{i-j}_{\mathrm{u}}(cool) = \mathsf{on}$, for $j \in 0..k-1$.*

Proof. Similar to the proof of Lemma 1

□

The following Lemma is a variant of Proposition 1. Differently from Proposition 1, the system behaves correctly until the attack A_n defined in Example 6 is alive. Indeed, the attack hides the real value of the temperature and hence an alarm can be sent only when the attack is terminated.

Lemma 3. *Let Sys be the system defined in Example 2 and A_n be the attack of Example 6. If $Sys \parallel A_n \xrightarrow{t}$, for some trace t containing less than n* idle *actions, then t contains only actions in $\{\tau, \mathsf{idle}\}$.*

Proof. Let $Sys \parallel A_n \xrightarrow{t}$ for some trace t containing less than n idle-actions. By Lemma 2 and since $\delta = 0.4$ the temperature of the system will remain always in the real interval $[0, 13.5]$. Thus, the invariant is never violated and the trace t does not contain any deadlock-action.

Let us demonstrate that the trace cannot contain outputs on the channel $alarm$ (the only other possible action). When the process IDS tests the temperature, the coolant has been active for (at least) 5 time slots. Since this test occurs at most in the n-th time slot, then, by Lemma 2, the coolant was active for exactly 5 slots. As a consequence, the state variable $temp$ will be in the real interval $(9.9 + 2 - 5 * (1 + \delta), 11.1 + 2 + \delta - 5 * (1 - \delta)] = (4.9, 10.5]$. Since the process IDS tests the temperature at most in the n-th time slot, the attack is alive and so IDS will sense a temperature less than or equal to 8.6 and it will not send an $alarm$. □

Lemma 4 is a variant of Proposition 2 where the CPS Sys runs in parallel with the attack A_n of Example 6. Similarly to Lemma 2, due to the presence of the attack, the temperature is 2 degrees higher with respect to the system in isolation.

Lemma 4. *Let Sys be the system defined in Example 2 and A_n be the attack of Example 6. Then, for any execution trace of $Sys \parallel A_n \xrightarrow{t}$ where t contains at most $n - 1$* idle *action we have the following:*

- *if either process $Ctrl$ or process IDS sense a temperature above 10 then the value of the state variable $temp$ ranges over $(9.9 + 2, 11.1 + 2 + \delta]$;*
- *when the process IDS tests the temperature the value of the state variable $temp$ ranges over $(9.9 + 2 - 5 * (1 + \delta), 11.1 + 2 + \delta - 5 * (1 - \delta)]$.*

Proof. Let us prove the two statements separately.

- Since $\epsilon = 0.1$, if either $Ctrl$ or IDS sense a temperature above 10 before the n-th time slot (and hence the attack is alive and Sys turns on the cooling) then the value of the state variable $temp$ is greater than $9.9 + 2 = 11.9$. By Lemma 2 the value of the state variable $temp$ is less or equal than $11.1 + 2 + \delta$. Therefore, if $Ctrl$ senses a temperature above 10 before the n-th slot, then the value of the state variable $temp$ is in $(9.9 + 2, 11.1 + 2 + \delta]$.
- Furthermore, if the process IDS tests the temperature before the n-th time slot, it means the the coolant it is active for at least 5 idle actions. By Lemma 2, the coolant it was active for exactly 5 idle actions and so, the state variable $temp$ will be in the real interval $(9.9 + 2 - 5 * (1 + \delta), 11.1 + 2 + \delta - 5 * (1 - \delta)]$.

□

Now, everything is in place to prove Proposition 6.

Proof of Proposition 6. Let us proceed by case analysis.

- Let $0 \leq n \leq 8$. In the proof of Proposition 4, we remarked that the system Sys in isolation may sense a temperature greater than 10 only after 8 idle-actions, i.e., in the 9-th time slot. However, the life of the attack is $n \leq 8$, and in the 9-th time slot the attack is already terminated. As a consequence, starting from the 9-th time slot the system will correctly sense the temperature and it will correctly activate the cooling system.
- Let $n = 9$. The maximum value that may be reached by the state variable $temp$ after 8 idle-actions, i.e., in the 9-th time slot, is $8 * (1 + \delta) = 8 * 1.4 = 11.2$. However, since in the 9-th time slot the attack is still alive, the process $Ctrl$ will sense a temperature below 10 and the system will move to the next time slot. Then, in the 10-th time slot, when the attack is already terminated, the maximum temperature the system may reach is $11.2 + (1 + \delta) = 12.6$ degrees. Thus, the process $Ctrl$ will sense a temperature greater than 10, activating the cooling system. As a consequence, after 5 time units of cooling, in the 15-th time slot, the value of the state variable $temp$ will be at most $12.6 - 5 * (1 - \delta) = 12.6 - 5 * 0.6 = 9.6$, which will be sensed by process IDS as at most 9.7 (sensor error $\epsilon = 0.1$) and no alarm will be turned on. Moreover, the invariant is obviously always preserved. As the attack made is already terminated, from this point in time on, the system will behave correctly with neither deadlocks or alarms.
- Let $10 \leq n \leq 15$. By Proposition 1 we know that the system Sys can only perform τ- and idle-actions. Thus, in order to prove that $Sys \parallel A_n \simeq_{[16,n+6]} Sys$, it is enough to show the following three facts:
  - the system $Sys \parallel A_n$ does not deadlock;
  - the system $Sys \parallel A_n$ may not emit any alarm in the first 15 time slots;
  - there is a trace in which the system $Sys \parallel A_n$ emits an output on channel $alarm$, exactly after 15 idle-actions, in the 16-th time slot;
  - there is a trace in which the system $Sys \parallel A_n$ emits an output along channel $alarm$, exactly after $n+5$ idle-actions, in the $n+6$-th time interval,
  - the system $Sys \parallel A_n$ does not have any trace emitting an output along channel $alarm$ after the $n+6$-th time slot.

  Concerning the first fact, since $\delta = 0.4$, by Lemma 2 the temperature of the system will remain always in the real interval $[0, 13.5]$ during the attack. Thus, the invariant is never violated and the trace of the system under attack cannot contain any deadlock-action. Moreover, when the attack terminates, if the temperature is in $[0, 9.9]$, the system continuous his behaviour correctly as in isolation. Otherwise, since the temperature is at most 13.5, after a possible sequence of the cooling cycles, the temperature will reach a value in $[0, 9.9]$ and the system will continue its behaviour correctly, as if in isolation.

  Concerning the second fact, in the previous case ($n = 9$), we have proven that in the 10-th time slot the temperature may reach at most the value 12.6. Thus, even in the presence of the attack, the controller activates the cooling, and after 5 time units of cooling, in the 15-th time slot, the temperature will reach at most 9.6 and no alarm will be sent.

  Concerning the third fact, according to the proof for $n = 9$, in the 9-th time slot the system may reach at most a temperature of 11.2 degrees. Thus, according to the evolution law, in the 10-th time slot the system may reach at most a temperature of $11.2 + 1 - 0.1 = 12.1$ degrees, with $0.1 \in [-\delta, +\delta]$. Now, since in the 10-th time slot the attack is still alive, and the sensor error is $\epsilon = 0.1$, there is an execution of the system in which the process $Ctrl$ senses a temperature of 10 degrees (hence, not above 10) moving the system into the next time slot. Thus, in the 11-th time slot, the system may reach a temperature of $12.1 + (1 + \delta) = 13.5$ degrees. In this case, independently of whether the attack is dead or alive, the process $Ctrl$ will sense a temperature above 10 and it will activate the cooling. As a consequence, after 5 time units of cooling, i.e., in the 16-th time slot, the value of the state variable $temp$ may be at most $13.5 - 5 * (1 - \delta) = 10.5$. In this case, as the attack is already terminated, the process IDS will sense a temperature above 10 and it will turn on the alarm. Note that we have chosen an execution trace in order to induce an alarm as soon as possible. It is not possible to produce an alarm in a previous time slot.

  Concerning the fourth fact, first we prove that there is a trace of the system $Sys \parallel A_n$ in which the state variable $temp$ reaches the value 12 in the time slot n. Since the attack is alive and $\epsilon = 0.1$, when the variable $temp$ is equal to 12 the sensed temperature is in $[9.9, 10.1]$. Thus, it suffices to increase the temperature of $\frac{12}{n}$ (since $10 \leq n \leq 15$ $\frac{12}{n}$ is an admissible increasing value). As a consequence, we can suppose that in the n-th time slot there is a trace in which the value of the variable $temp$ is 12. Since in that time slot the attack is alive and $\epsilon = 0.1$, the sensed temperature lays in the interval $[9.9, 10.1]$. Let us focus on the trace in which the sensed temperature is 10 and the cooling system is not activated. In this case, in the $n+1$-th time slot the system may reach a temperature of $12 + 1 + \delta = 13.4$ degrees and the process $Ctrl$ will sense a temperature above 10 activating the cooling system. As a consequence, after further 5 time units of cooling, i.e. in the $n+6$-th time slot, the value of the state variable $temp$ may be $13.5 - 5 * (1 - \delta) = 10.4$ and the process IDS will fire an alarm.

  Concerning the fifth fact, we can distinguish two cases: (i) there was no alarm before the $n + 6$-th time slot, or (ii) there was an alarm in the m-th time interval, with $m \in [16, n + 6]$. Let us start with the case (i). Since the attack terminated more than 5 time slot ago, and since no alarm fired in the meanwhile, it follows that the possible

activation of the refrigerant drove the temperature under the desired thresholds. So, after the $n+6$-th time slot the system will keep behaving correctly, as if in isolation. Consider now the case (ii), i.e., an alarm occurred in the m-th time interval, with $m \in [16, n+6]$. By Lemma 2, an alarm occurring in the m-th time interval follows from an activation of the refrigerant that occurred in the $(m-5)$-th time interval with a temperature that was at most 13.5. Thus, the alarm that occurred in the m-th interval is due to a temperature that must be at most $13.5 - 5 * (1-\delta) = 13.5 - 5 * 0.6 = 10.5$. Therefore, after further 5 time units of refrigerant, the temperature will be at most $10.5 - 5 * (1-\delta) = 10.5 - 5 * 0.6 = 7.5$. As a consequence, after the $n+6$-th time slot, no alarm will be fired and the system will behave correctly, as if in isolation.

- Let $n \geq 16$. In order to prove that $Sys \parallel A_n \simeq_{[n+1,n+6]} Sys$, it is enough to show that:
  - the system $Sys \parallel A_n$ does not deadlock;
  - the system $Sys \parallel A_n$ may not emit any output in the first n time slots;
  - there is a trace in which the system $Sys \parallel A_n$ emits an output on channel $alarm$ in the $(n+1)$-th time slot;
  - there is a trace in which the system $Sys \parallel A_n$ emits an output along channel $alarm$ in the $n+6$-th time slot;
  - the system $Sys \parallel A_n$ does not have any execution trace emitting an output along channel $alarm$ after the $n+6$-th time slot.

  Concerning the first fact, the proof is analogous to that of the first fact of the case when $10 \leq n \leq 15$.
  Concerning the second fact, when the process $Check$ tests the temperature (provided by the attack), by Lemma 4 (a straightforward variant of Proposition 2), it follows that the state variable $temp$ can range in the interval $(4.9, 10.5]$. As $\epsilon = 0.1$ and the attack is alive, the sensed temperature is at most 8.6 and no alarm will be emitted.
  Concerning the third fact, firstly we prove that for all time slot m, with $9 \leq m < n$, there is a trace of the system $Sys \parallel A_n$ in which the state variable $temp$ reaches the values 12 in the time slot m. Since the attack is alive and $\epsilon = 0.1$, when the variable $temp$ is equal to 12 the sensed temperature is in $[9.9, 10.1]$.
  The fastest trace reaching the temperature of 12 degree requires $\lceil \frac{12}{1+\delta} \rceil = \lceil \frac{12}{1.4} \rceil = 9$ time units, whereas the slowest one $\lceil \frac{12}{1-\delta} \rceil = \lceil \frac{12}{0.6} \rceil = 20$ time units. Thus, for any time slot m, with $9 < m \leq 21$, there is a trace of the system where the value of the state variable $temp$ is 12. Now, for any of those time slots m there is a trace in which the state variable $temp$ is equal to 12 in all time slots $m + 10i < n$, with $i \in \mathbb{N}$. As already said, when the variable $temp$ is equal to 12 the sensed temperature varies lays in the interval $[9.9, 10.1]$ and the cooling might be activated. Thus, there is a trace in which the cooling system is activated. We can always assume that during the cooling the temperature decreases of $1+\delta$ degrees per time unit, reaching at the end of the cooling cycle the value of 5. Thus, the trace may continue with 5 time slots in which the variable $temp$ is increased of $1+\delta$ degrees per time unit; reaching again the value 12. Thus, for all time slots m, with $9 < m < n$, there is a trace of the system $Sys \parallel A_n$ in which the state variable $temp$ is 12 in m.
  As a consequence, we can suppose that in the $n-5$-th time slot there is a trace in which the value of the variable $temp$ is 12. Since in that time slot the attack is alive and $\epsilon = 0.1$, the sensed temperature lays in the interval $[9.9, 10.1]$. Let us focus on the trace in which the sensed temperature is 10 and the cooling system is not activated. In this case, in the $n-4$-th time slot the system may reach a temperature of $12 + (1+\delta) = 13.4$ degrees and the process $Ctrl$ will sense a temperature above 10 activating the cooling system. As a consequence, after further 5 time units of cooling, i.e., in the $n+1$-th time slot, the value of the state variable $temp$ may be $13.5 - 5 * (1-\delta) = 10.4$ and the alarm will be fired.
  Concerning the fourth fact, this can be proved as the previous one by considering that in the n-th time interval there is a trace in which the value of the variable $temp$ is 12. Thus, in the $n+1$-th time interval the system may reach a temperature of $12 + 1 + \delta = 13.4$ degrees and the process $Ctrl$ will sense a temperature above 10 activating the cooling system. As a consequence, after further 5 time units of cooling, i.e., in the $n+6$-th time interval, the value of the state variable $temp$ may reach $13.5 - 5 * (1-\delta) = 10.4$ and the alarm will be fired.
  Finally, the fifth fact can be proved as the fifth fact of the case when $10 \leq n \leq 15$.

□

Proof of Proposition 7. Firstly, let us show that $Sys \parallel A \simeq Sys$. We recall that by Proposition 1, the system Sys never deadlocks and never sends alarm: it can only perform τ-actions and idle-actions. Let us examine the possible execution traces of $Sys \parallel A$. Notice that in the first 10 time slots the attack A only reads the sensor s_t and hence it does not affect Sys. In the 11-th time slot, if the attack senses a temperature equal to 9.1, then, since $\epsilon = 0.1$, the state variable $temp$ is in the interval $[9, 9.2]$ and hence the process $Ctrl$ does not activate the cooling system and moves to the next time slot. Since $\delta = 0.4$, in the 12-th time slot, the state variable $temp$ may be in the real interval $[9.6, 10.6]$ and hence, since $\epsilon = 0.1$, the sensed temperature will be in the interval $[9.5, 10.7]$. Now, there are two cases:

- If the sensed temperature is above 10, then, after 5 time units of cooling, in the 17-th time slot, the variable $temp$ will be in the real interval $[2.6, 6.6]$. At this point in time the attack is already terminated and the system will proceed correctly, without deadlocks or alarms.
- If the sensed temperature is not above 10, then the system moves in the 13-rd time slot where the variable $temp$ will move in the real interval $[10.2, 12]$. Now, since $\epsilon = 0.1$ the sensed temperature is surely above 10

and the cooling is activated. Thus, after 5 time units of cooling, in the 18-th time slot, the variable $temp$ will drop in the real interval $[3.2, 9]$. Now, since the attack is already terminated, the system will proceed correctly, without deadlocks or attacks.

Let us show now that $Sys' \parallel A \simeq_{[16,\infty]} Sys'$. By a straightforward variant of Proposition 1, it easy to show that the system Sys' never deadlocks and never sends alarm. Actually, systems Sys and Sys' are trace equivalent. Thus, in order to show our system equality it is enough to show the following facts:

- the system $Sys' \parallel A$ may not deadlock for the first 15 time slots;
- the system $Sys' \parallel A$ may not emit any output for the first 15 time slots;
- there is a trace in which the system $Sys' \parallel A$ emits an output on channel $alarm$, in the 16-th time slot, and then, after 2 time slots, the system deadlocks.

The proofs of the first two facts are similar to that seen in the previous proofs. Concerning the third fact, if the attack senses a temperature equal to 9.1 in the 11-th time slot, then the same temperature might be sensed by process $Ctrl'$. In that case, $temp \in [9, 9.2]$ and the controller activates the cooling system. Now, suppose that during the cooling the temperature decreases of $1 + \delta = 1.4$ per time unit. Then, in the 16-th time slot the temperature will be in the interval $[9 - 5 * 1.4, 9.2 - 5 * 1.4] = [2, 2.2]$. However, after 5 time units of cooling, in the 16-th time slot, the controller will check again the temperature and it will get the value 10.5 provided by the attacker. As a consequence, the process IDS will emit an alarm in the 16-th time slot. Notice that no alarm can be sent in previous time slots. Moreover the process IDS says to $Ctrl$ to keep cooling. Since the temperature is in $[2, 2.2]$, after further 2 time units of cooling the temperature is in $[2 - 1.4 * 2, 2.2 - 0.6 * 2] = [-0.8, 1]$, and hence the invariant could be violated, as the temperature could be below zero. □

3. Proofs of § 4

In order to prove Proposition 9 we need a couple of lemmas.

Lemma 5 is a variant of Lemma 1. Here the behaviour of Sys is parametric on the uncertainty.

Lemma 5. *Let Sys be the system defined in Example 2, and $0.4 < \gamma \leq \frac{19}{30}$. Let*

$$Sys\{^{\gamma}/_{\delta}\} = Sys_1 \xrightarrow{t_1} \xrightarrow{\mathsf{idle}} Sys_2 \xrightarrow{t_2} \xrightarrow{\mathsf{idle}} \cdots \xrightarrow{t_{n-1}} \xrightarrow{\mathsf{idle}} Sys_n$$

such that the traces t_j contain no idle*-actions, for any $j \in 1..n-1$, and for any $i \in 1..n$ $Sys_i = E_i \bowtie P_i$ with $E_i = \langle \xi^i_{\mathrm{x}}, \xi^i_{\mathrm{u}}, \gamma, evol, \epsilon, meas, inv \rangle$. Then, for any $i \in 1..n-1$ we have the following:*

- *if $\xi^i_{\mathrm{u}}(cool) = \mathsf{off}$ then $\xi^i_{\mathrm{x}}(temp) \in [0, 11.1 + \gamma]$;*
- *if $\xi^i_{\mathrm{u}}(cool) = \mathsf{off}$ and $\xi^i_{\mathrm{x}}(temp) \in (10.1, 11.1 + \gamma]$ then $\xi^{i+1}_{\mathrm{u}}(cool) = \mathsf{on}$;*
- *if $\xi^i_{\mathrm{u}}(cool) = \mathsf{on}$ then $\xi^i_{\mathrm{x}}(temp) \in (9.9 - k * (1 + \gamma), 11.1 + \gamma - k * (1 - \gamma)]$, for some $k \in 1..5$, such that $\xi^{i-k}_{\mathrm{u}}(cool) = \mathsf{off}$ and $\xi^{i-j}_{\mathrm{u}}(cool) = \mathsf{on}$, for $j \in 0..k-1$.*

Proof. Similar to the proof of Lemma 1 □

The following Lemma is a variant of Proposition 1.

Lemma 6. *Let Sys be the system defined in Example 2 and γ such that $0.4 < \gamma \leq \frac{19}{30}$. If $Sys\{^{\gamma}/_{\delta}\} \xrightarrow{t} Sys'$, for some $t = \alpha_1 \ldots \alpha_n$, then $\alpha_i \in \{\tau, \mathsf{idle}\}$, for any $i \in 1..n$.*

Proof. By Lemma 5 the temperature is always in $[0, 11.1+\gamma]$ and so, since $\gamma \leq \frac{19}{30}$ the system never deadlocks. Moreover, after 5 idle action of coolant the state variable $temp$ is in $(9.9-5*(1+\gamma), 11.1+\gamma-5*(1-\gamma)] = (4.9-5\gamma\,,\ 6.1+6\gamma]$. Since $\epsilon = 0.1$, the value detected from the sensor will be in the interval $(4.8 - 5\gamma\,,\ 6.2 + 6\gamma]$. Thus, the temperature sensed by IDS will be at most $6.2 + 6\gamma \leq 6.2 + 6 * \frac{19}{30} = \frac{6.2*5+19}{5} = 10$. Hence no alarm is fired. □

The following Lemma is a variant of Proposition 2. Here the behaviour of Sys is parametric on the uncertainty.

Lemma 7. *Let Sys be the system defined in Example 2 and γ such that $0.4 < \gamma \leq \frac{19}{30}$. Then, for any execution trace of $Sys\{^{\gamma}/_{\delta}\}$ we have the following:*

- *if either process $Ctrl$ or process IDS senses a temperature above 10 then the value of the state variable $temp$ ranges over $(9.9, 11.1 + \gamma]$;*
- *when the process IDS tests the temperature the value of the state variable $temp$ ranges over $(9.9 - 5 * (1 + \gamma), 11.1 + \gamma - 5 * (1 - \gamma)]$.*

Proof. As to the first statement, since $\epsilon = 0.1$, if either process $Ctrl$ or process IDS senses a temperature above 10 then the value of the state variable $temp$ is above 9.9. By Lemma 5 the state variable $temp$ is less or equal than $11.1 + \gamma$. Therefore, *if either process $Ctrl$ or process IDS sense* a temperature above 10 then the value of the state variable $temp$ is in $(9.9, 11.1 + \gamma]$.

Let us prove now the second statement. When the process IDS tests the temperature then the coolant has been active for 5 idle actions. By Lemma 5 the state variable $temp$ ranges over $(9.9 - 5 * (1 + \gamma), 11.1 + \gamma - 5 * (1 - \gamma)]$. □

Finally, everything is in place to prove Proposition 9.

Proof of Proposition 9. For (1) we have to show that the two systems Sys and $Sys\{^{\gamma}/_{\delta}\}$ are trace equivalent, for $\gamma \in (\frac{12}{30}, \frac{19}{30})$. By Proposition 1 it is enough to show that the system $Sys\{^{\gamma}/_{\delta}\}$ never deadlock and never sends alarm. But this obviously holds by Lemma 6. This is enough to derive that $Sys \simeq Sys\{^{\gamma}/_{\delta}\}$, for $\gamma \in (\frac{12}{30}, \frac{19}{30})$.

For (2), we have to show that the two systems Sys and $Sys\{^{\gamma}/_{\delta}\}$ are not trace equivalent, for $\gamma > \frac{19}{30}$. By Proposition 1 it is enough to show that the system $Sys\{^{\gamma}/_{\delta}\}$ has a trace which either sends an alarm or deadlocks. We can easily build up a trace for $Sys\{^{\gamma}/_{\delta}\}$ in which, after 10 idle-actions, in the 11-th time slot, the value of the state variable

$temp$ is 10.1. In fact, it is enough to increase the temperature of 1.01 degrees for the first 10 rounds. Notice that this is an admissible value since, $1.01 \in [1-\gamma, 1+\gamma]$, for any $\gamma > \frac{19}{30}$. Being 10.1 the value of the state variable $temp$, there is an execution trace in which the sensed temperature is 10 (recall that $\epsilon = 0.1$) and hence the cooling system is not activated. However, in the following time slot, i.e., the 12-th time slot, the temperature may reach at most the value $10.1+1+\gamma$. Now, if $10.1 + 1 + \gamma > 50$ then the system deadlocks. Otherwise, the controller activates the cooling system, and after 5 time units of cooling, in the 17-th time slot, the variable $temp$ will be at most $11.1 + \gamma - 5(1 - \gamma) = 6.1 + 6\gamma$. Since $\epsilon = 0.1$, the sensed temperature would be in the real interval $[6.0+6\gamma, 6.2+6\gamma]$. Thus, there is an execution trace in which the sensed temperature is $6.2 + 6\gamma$, which will be greater than 10 being $\gamma > \frac{19}{30}$. Thus, the alarm will be transmitted.

This is enough to derive that $Sys \not\simeq Sys\{^{\gamma}/_{\delta}\}$, for $\gamma > \frac{19}{30}$. □

Proof of Proposition 10. Let us prove the first sub-result. From Proposition 6 we know that $Sys \parallel A_{10} \simeq_{15} Sys$. In particular we showed that the system $Sys \parallel A_{10}$ has an execution trace which fires an alarm in the 16-th time slot, and which cannot be matched by Sys. As a consequence, $Sys \parallel A_{10} \not\simeq Sys$. Furthermore, in Proposition 9 we showed that the system Sys is $\frac{7}{30}$-tolerant. Formally, for $\gamma \in (\frac{12}{30}, \frac{19}{30})$ we have $Sys \simeq Sys\{^{\gamma}/_{\delta}\}$. Thus, for $\gamma \in (\frac{12}{30}, \frac{19}{30})$ we have

$$Sys \parallel A_{10} \not\simeq Sys\{^{\gamma}/_{\delta}\} \ .$$

Let us prove the second sub-result. That is,

$$Sys \parallel A_{10} \simeq Sys\{^{\gamma}/_{\delta}\}$$

for $\gamma > \frac{19}{30}$. Formally, we have to demonstrate that whenever $Sys \parallel A_{10} \xrightarrow{t}$, for some trace t, then $Sys\{^{\gamma}/_{\delta}\} \stackrel{\hat{t}}{\Longrightarrow}$ as well. Let us do a case analysis on the structure of the trace t. We distinguish three possible cases.

- The trace t contains deadlock-actions. This case cannot occur. In fact, in Example 8 we showed that the system $Sys \mid A_{10}$ may not deadlock in the first 15 time slots. Furthermore, in 11-th time slot the attack A_{10} is terminated and it does not have any influence on the system Sys. Thus, after an eventual alarm the system will recover its correct behaviour with no deadlocks.
- The trace t contains only τ- and idle-actions. If the system $Sys \parallel A_{10}$ has a trace t which contains only τ- and idle-actions then, by Proposition 1, the system Sys in isolation must have a similar trace with the same number of idle-actions. By an application of Proposition 8, as $\delta < \gamma$, any trace of Sys can be simulated by $Sys\{^{\gamma}/_{\delta}\}$. As a consequence, $Sys\{^{\gamma}/_{\delta}\} \stackrel{\hat{t}}{\Longrightarrow}$.
- The trace t contains only τ-, idle- and $\overline{alarm}$-actions. Firstly, notice that the system $Sys \parallel A_{10}$ may produce *only one* output on channel $alarm$, in the 16-th time slot. After that, the trace will have only τ- and idle-actions. In fact, in Proposition 6 we provided a trace in which, in the 16-th time slot, the state variable $temp$ reaches the value 10.5 and an output on channel $alarm$ is emitted. This is the maximum possible value for variable $temp$ in that point in time. After the transmission of the $alarm$, the system $Sys \parallel A_{10}$ activates the cooling for the following 5 time slots. Thus, in the 21-th time slot, the temperature will be at most $10.5 - 5 * (1 - \delta) = 10.5 - 5 * (0.6) = 7.5$, and no alarm is fired. From that time on, since the attack A_{10} terminated is life in the 10-th time slot, no other alarms will be fired. Summarising $Sys \parallel A_{10}$ fires only one alarm in the 16-th time slot.
  It remains to prove is that for any $\gamma > \frac{19}{30}$ the system $Sys\{^{\gamma}/_{\delta}\}$ has a trace which performs an output on channel $alarm$, exactly in the 16-th time slot. By Proposition 8 it is enough to prove that such trace exists for $\frac{19}{30} < \gamma < \frac{20}{30}$ to derive that it exists also for $\gamma \geq \frac{20}{30}$.
  We can easily build up a trace for $Sys\{^{\gamma}/_{\delta}\}$ in which, after 9 idle-actions, in the 10-th time slot, the value of the state variable $temp$ is 10.1. In fact, it is enough to increase the temperature of $\frac{10.1}{9}$ degrees for the first 9 rounds. Notice that this is an admissible value since, $\frac{10.1}{9} = 1.1\overline{2} \in [1 - \gamma, 1 + \gamma]$, for any $\frac{19}{30} < \gamma < \frac{20}{30}$. Being 10.1 the value of the state variable $temp$, there is an execution trace in which the sensed temperature is 10 (recall that $\epsilon = 0.1$) and hence the cooling system is not activated. However, in the following time slot, i.e. the 11-th time slot, the temperature may reach at most the value $10.1 + 1 + \gamma$, imposing the activation of the cooling system. After 5 time units of cooling, in the 16-th time slot, the variable $temp$ will be at most $11.1 + \gamma - 5(1 - \gamma) = 6.1 + 6\gamma$. Since $\epsilon = 0.1$, the sensed temperature would be in the real interval $[6.0+6\gamma, 6.2+6\gamma]$. Thus, there is an execution trace in which the sensed temperature is $6.2+6\gamma$, which will be greater than 10 being $\gamma > \frac{19}{30}$. Thus, the alarm will be transmitted, in the 16-th time slot, as required. After the transmission on channel $alarm$, the system $Sys\{^{\gamma}/_{\delta}\}$ activates the cooling for the following 5 time slots. As a consequence, in 21-th time slot, the temperature will be at most $6.1 + 6\gamma - 5 * (1 + \gamma) = 1.1 + \gamma$. Since we assumed $\frac{19}{30} < \gamma < \frac{20}{30}$ the temperature will be well below 10 and no alarm will be sent. From that time on, since $Sys\{^{\gamma}/_{\delta}\}$ can mimic all traces of Sys, we can always choose a trace which does not fire the alarm any more.

This is enough to derive that:

$$Sys \parallel A_{10} \simeq Sys\{^{\gamma}/_{\delta}\} \ .$$

□

Proof of Proposition 11. Let us prove the first sub-result. As argued in Example 4 we know that $Sys \parallel A \simeq_{36} Sys$ because $Sys \parallel A$ deadlocks in the 37-th time slot. Furthermore, for $\gamma \in (0.4, 0.98)$ the system $Sys\{^{\gamma}/_{\delta}\}$ cannot deadlock. Indeed, by an easy extension of the first statement of Lemma 5, for $\gamma \in (0.4, 0.98)$ the coolant can be activated only when the temperature is in $(9.9, 10.1+(1+\gamma)]$. Therefore, since $\gamma \in (0.4, 0.98)$, the temperature will always

be below 50. Moreover, since $\epsilon = 0.1$, when the coolant is activated we know that the temperature is above 9.9. Thus, after 5 time units of cooling the temperature is at least $9.9-5*(1+\gamma) = 4.9-5*\gamma$. Since $\gamma \in (0.4, 0.98)$, the state variable *temp* cannot drop below 0. As a consequence, the invariant is always satisfied and the system never deadlocks. This implies that

$$Sys \parallel A \not\simeq Sys\{^{\gamma}/_{\delta}\} \ .$$

Let us prove the second sub-result. That is,

$$Sys \parallel A \simeq Sys\{^{\gamma}/_{\delta}\}$$

for $\gamma > 0.98$. We demonstrate that whenever $Sys \parallel A \xrightarrow{t}$, for some trace t, then $Sys\{^{\gamma}/_{\delta}\} \stackrel{\hat{t}}{\Longrightarrow}$ as well. We will proceed by case analysis on the kind of actions contained in t. We distinguish three possible cases.

- The trace t contains only τ-, idle- and $\overline{alarm}$-actions. This case cannot occur. In fact as discussed in Example 4, the process *Ctrl* never activates the *Cooling* component (and hence also the *IDS* component which is the only one which could send an alarm) since it will always detect a temperature below 10.
- The trace t contains only τ- and idle-actions.
  If the system $Sys \parallel A$ has a trace t which contains only τ- and idle-actions then, by Proposition 1, the system Sys in isolation must have a similar trace with the same number of idle-actions. By an application of Proposition 8, as $\delta < \gamma$, any trace of Sys can be simulated by $Sys\{^{\gamma}/_{\delta}\}$. As a consequence, $Sys\{^{\gamma}/_{\delta}\} \stackrel{\hat{t}}{\Longrightarrow}$.
- The trace t contains contains only τ-, idle- and deadlock-actions.
  As discussed in Example 4, the system $Sys \parallel A$ will eventually deadlock. This is because the process *Ctrl* never activates the coolant (and hence the *IDS* never sends the alarm) since it will always detect a temperature below 10. Thus, the fastest execution deadlocks after $\lceil \frac{50}{1+\delta} \rceil = \lceil \frac{50}{1.4} \rceil = 36$ idle-actions, while the slowest one after $\lceil \frac{50}{1-\delta} \rceil = \lceil \frac{50}{0.6} \rceil = 84$ idle-actions.
  As a consequence, if $Sys \parallel A \xrightarrow{t}$, contains a deadlock-actions, then it will be after a number of idle-actions which is between 36 and 84. Hence we have to prove that for any $\gamma > 0.98$ and $n \in [36, 84]$, the system $Sys\{^{\gamma}/_{\delta}\}$ has a trace containing only τ and idle-actions which deadlocks after n idle actions. By Proposition 8 it is enough to show that such trace exists for $0.98 < \gamma < 1$. Indeed, if that trace exists for $0.98 < \gamma < 1$, then it exists also for $\gamma \geq 1$.
  Given $n \in [36, 84]$ we construct a trace deadlocking after n idle actions.
  As showed in the proof of Proposition 4 (second case), we can construct a trace for Sys (and hence by monotonicity for $Sys\{^{\gamma}/_{\delta}\}$) which activates the cooling exactly in the $(n-14)$-th time slot. Again, by monotonicity and by Proposition 2, the system $Sys\{^{\gamma}/_{\delta}\}$, after 5 idle actions of cooling, reaches in the $(n-9)$-th time slot a temperature $v \in (2.9, 8.5]$ for which the cooling is switched off and no alarm is fired.
  Let k such that $0 < k < \gamma - 0.98$ (k exists since $\gamma > 0.98$) we consider an increment of temperature of $\frac{9.9+5*k-v}{5}$ for 5 idle actions in such a way that at the $(n-4)$-th time slot the temperature becomes $9.9 + 5 * k$. We note that $\frac{9.9+5*k-v}{5}$ is an admissible increment since $\frac{9.9+5*k-v}{5} > \frac{9.9+0-8.5}{5} > 0.02 > 1-\gamma$ and $\frac{9.9+5*k-v}{5} < \frac{9.9+5*(\gamma-0.98)-2.9}{5} < \frac{9.9+5*(1-0.98)-2.9}{5} < \frac{9.9+0.2-2.9}{5} < 1.98 < 1+\gamma$.
  Now, since in $(n-4)$-th time slot the temperature is $9.9 + 5 * k$ and since $k > 0$, the cooling is switched on. Finally, we consider a decrement of temperature of $1+\gamma$ for 5 idle actions. Therefore at $(n+1)$-th time slot, the temperature is equal to $9.9 + 5 * k - 5 * (1 + \gamma) = 4.9 + 5 * (k - \gamma)$ and, since $0 < k < \gamma - 0.98$, we have that $4.9 + 5 * (k - \gamma) < 4.9 + 5 * (-0.98) = 0$. Hence the temperature is less than 0 and so the system deadlocks at the $(n+1)$-th time slot.
  Summarising, for any $\gamma > 0.98$ and $n \in [36, 84]$ the system $Sys\{^{\gamma}/_{\delta}\}$ has a trace containing only τ and idle actions which deadlocks after n idle actions.

This is enough to derive that:

$$Sys \parallel A \simeq Sys\{^{\gamma}/_{\delta}\} \ .$$

□